



# Red Hat build of MicroShift 4.15

## API reference

MicroShift API reference



## Red Hat build of MicroShift 4.15 API reference

---

MicroShift API reference

## Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

This document provides details about the REST API functions in MicroShift.

# Table of Contents

<b>CHAPTER 1. UNDERSTANDING API TIERS</b> .....	<b>50</b>
1.1. API TIERS	50
API tier 1	50
API tier 2	50
API tier 3	50
API tier 4	50
1.2. MAPPING API TIERS TO API GROUPS	50
1.2.1. Support for Kubernetes API groups	51
1.2.2. Support for OpenShift API groups	51
1.3. API DEPRECATION POLICY	51
1.3.1. Deprecating parts of the API	51
1.3.2. Deprecating CLI elements	52
1.3.3. Deprecating an entire component	53
<b>CHAPTER 2. UNDERSTANDING API COMPATIBILITY GUIDELINES</b> .....	<b>54</b>
2.1. API COMPATIBILITY GUIDELINES	54
2.2. API COMPATIBILITY EXCEPTIONS	55
Functional defaults between an upgraded cluster and a new installation	55
Usage of API fields that have the prefix "unsupported" or undocumented annotations	55
API availability per product installation topology	55
2.3. API COMPATIBILITY COMMON TERMINOLOGY	55
2.3.1. Application Programming Interface (API)	55
2.3.2. Application Operating Environment (AOE)	55
2.3.3. Compatibility in a virtualized environment	56
2.3.4. Compatibility in a cloud environment	56
2.3.5. Major, minor, and z-stream releases	56
2.3.6. Extended user support (EUS)	56
2.3.7. Developer Preview	56
2.3.8. Technology Preview	57
<b>CHAPTER 3. API INDEX</b> .....	<b>58</b>
<b>CHAPTER 4. API OBJECT REFERENCE</b> .....	<b>61</b>
4.1. COMMON OBJECT REFERENCE	61
4.1.1. io.k8s.api.admissionregistration.v1.MutatingWebhookConfigurationList schema	61
Schema	61
4.1.2. io.k8s.api.admissionregistration.v1.ValidatingWebhookConfigurationList schema	62
Schema	62
4.1.3. io.k8s.api.apps.v1.ControllerRevisionList schema	63
Schema	63
4.1.4. io.k8s.api.apps.v1.DaemonSetList schema	64
Schema	64
4.1.5. io.k8s.api.apps.v1.DeploymentList schema	65
Schema	65
4.1.6. io.k8s.api.apps.v1.ReplicaSetList schema	66
Schema	66
4.1.7. io.k8s.api.apps.v1.StatefulSetList schema	66
Schema	67
4.1.8. io.k8s.api.autoscaling.v2.HorizontalPodAutoscalerList schema	67
Schema	68
4.1.9. io.k8s.api.batch.v1.CronJobList schema	68
Schema	68

---

4.1.10. io.k8s.api.batch.v1.JobList schema	69
Schema	69
4.1.11. io.k8s.api.certificates.v1.CertificateSigningRequestList schema	70
Schema	70
4.1.12. io.k8s.api.coordination.v1.LeaseList schema	71
Schema	71
4.1.13. io.k8s.api.core.v1.Affinity schema	72
Schema	72
4.1.14. io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource schema	73
Schema	73
4.1.15. io.k8s.api.core.v1.AzureDiskVolumeSource schema	74
Schema	74
4.1.16. io.k8s.api.core.v1.AzureFilePersistentVolumeSource schema	75
Schema	75
4.1.17. io.k8s.api.core.v1.AzureFileVolumeSource schema	76
Schema	76
4.1.18. io.k8s.api.core.v1.Capabilities schema	76
Schema	77
4.1.19. io.k8s.api.core.v1.CephFSPersistentVolumeSource schema	77
Schema	77
4.1.20. io.k8s.api.core.v1.CephFSVolumeSource schema	78
Schema	78
4.1.21. io.k8s.api.core.v1.CinderPersistentVolumeSource schema	79
Schema	79
4.1.22. io.k8s.api.core.v1.CinderVolumeSource schema	80
Schema	80
4.1.23. io.k8s.api.core.v1.ClaimSource schema	81
Schema	81
4.1.24. io.k8s.api.core.v1.ComponentStatusList schema	82
Schema	82
4.1.25. io.k8s.api.core.v1.ConfigMapEnvSource schema	83
Schema	83
4.1.26. io.k8s.api.core.v1.ConfigMapKeySelector schema	84
Schema	84
4.1.27. io.k8s.api.core.v1.ConfigMapList schema	84
Schema	84
4.1.28. io.k8s.api.core.v1.ConfigMapProjection schema	85
Schema	85
4.1.29. io.k8s.api.core.v1.ConfigMapVolumeSource schema	86
Schema	86
4.1.30. io.k8s.api.core.v1.Container schema	87
Schema	88
4.1.31. io.k8s.api.core.v1.ContainerPort schema	93
Schema	94
4.1.32. io.k8s.api.core.v1.ContainerResizePolicy schema	94
Schema	95
4.1.33. io.k8s.api.core.v1.CSIPersistentVolumeSource schema	95
Schema	95
4.1.34. io.k8s.api.core.v1.CSIVolumeSource schema	97
Schema	97
4.1.35. io.k8s.api.core.v1.DownwardAPIProjection schema	98
Schema	98
4.1.36. io.k8s.api.core.v1.DownwardAPIVolumeFile schema	98

Schema	99
4.1.37. io.k8s.api.core.v1.DownwardAPIVolumeSource schema	99
Schema	100
4.1.38. io.k8s.api.core.v1.EmptyDirVolumeSource schema	100
Schema	100
4.1.39. io.k8s.api.core.v1.EndpointsList schema	101
Schema	101
4.1.40. io.k8s.api.core.v1.EnvFromSource schema	102
Schema	102
4.1.41. io.k8s.api.core.v1.EnvVar schema	102
Schema	103
4.1.42. io.k8s.api.core.v1.EnvVarSource schema	103
Schema	103
4.1.43. io.k8s.api.core.v1.EphemeralContainer schema	104
Schema	104
4.1.44. io.k8s.api.core.v1.EphemeralVolumeSource schema	110
Schema	110
4.1.45. io.k8s.api.core.v1.EventList schema	111
Schema	111
4.1.46. io.k8s.api.core.v1.EventSource schema	112
Schema	112
4.1.47. io.k8s.api.core.v1.ExecAction schema	113
Schema	113
4.1.48. io.k8s.api.core.v1.FCVolumeSource schema	113
Schema	113
4.1.49. io.k8s.api.core.v1.FlexPersistentVolumeSource schema	114
Schema	114
4.1.50. io.k8s.api.core.v1.FlexVolumeSource schema	115
Schema	115
4.1.51. io.k8s.api.core.v1.FlockerVolumeSource schema	116
Schema	116
4.1.52. io.k8s.api.core.v1.GCEPersistentDiskVolumeSource schema	116
Schema	117
4.1.53. io.k8s.api.core.v1.GitRepoVolumeSource schema	117
Schema	118
4.1.54. io.k8s.api.core.v1.GlusterfsPersistentVolumeSource schema	118
Schema	118
4.1.55. io.k8s.api.core.v1.GlusterfsVolumeSource schema	119
Schema	119
4.1.56. io.k8s.api.core.v1.GRPCAction schema	120
Schema	120
4.1.57. io.k8s.api.core.v1.HostAlias schema	120
Schema	121
4.1.58. io.k8s.api.core.v1.HostPathVolumeSource schema	121
Schema	121
4.1.59. io.k8s.api.core.v1.HTTPGetAction schema	122
Schema	122
4.1.60. io.k8s.api.core.v1.HTTPHeader schema	123
Schema	123
4.1.61. io.k8s.api.core.v1.ISCSIPersistentVolumeSource schema	123
Schema	124
4.1.62. io.k8s.api.core.v1.ISCSIVolumeSource schema	125
Schema	125

4.1.63. io.k8s.api.core.v1.KeyToPath schema	127
Schema	127
4.1.64. io.k8s.api.core.v1.Lifecycle schema	127
Schema	128
4.1.65. io.k8s.api.core.v1.LifecycleHandler schema	128
Schema	129
4.1.66. io.k8s.api.core.v1.LimitRangeList schema	129
Schema	129
4.1.67. io.k8s.api.core.v1.LocalObjectReference schema	130
Schema	130
4.1.68. io.k8s.api.core.v1.LocalVolumeSource schema	130
Schema	131
4.1.69. io.k8s.api.core.v1.NamespaceList schema	131
Schema	131
4.1.70. io.k8s.api.core.v1.NFSVolumeSource schema	132
Schema	132
4.1.71. io.k8s.api.core.v1.NodeAffinity schema	133
Schema	133
4.1.72. io.k8s.api.core.v1.NodeList schema	134
Schema	134
4.1.73. io.k8s.api.core.v1.NodeSelector schema	134
Schema	135
4.1.74. io.k8s.api.core.v1.NodeSelectorRequirement schema	135
Schema	135
4.1.75. io.k8s.api.core.v1.NodeSelectorTerm schema	136
Schema	136
4.1.76. io.k8s.api.core.v1.ObjectFieldSelector schema	136
Schema	136
4.1.77. io.k8s.api.core.v1.ObjectReference schema	137
Schema	137
4.1.78. io.k8s.api.core.v1.PersistentVolumeClaim schema	138
Schema	138
4.1.79. io.k8s.api.core.v1.PersistentVolumeClaimList schema	150
Schema	151
4.1.80. io.k8s.api.core.v1.PersistentVolumeClaimSpec schema	151
Schema	151
4.1.81. io.k8s.api.core.v1.PersistentVolumeClaimTemplate schema	154
Schema	155
4.1.82. io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource schema	155
Schema	155
4.1.83. io.k8s.api.core.v1.PersistentVolumeList schema	156
Schema	156
4.1.84. io.k8s.api.core.v1.PersistentVolumeSpec schema	156
Schema	157
4.1.85. io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource schema	161
Schema	162
4.1.86. io.k8s.api.core.v1.PodAffinity schema	162
Schema	162
4.1.87. io.k8s.api.core.v1.PodAffinityTerm schema	163
Schema	163
4.1.88. io.k8s.api.core.v1.PodAntiAffinity schema	164
Schema	164
4.1.89. io.k8s.api.core.v1.PodDNSConfig schema	165



Schema	165
4.1.90. io.k8s.api.core.v1.PodDNSConfigOption schema	166
Schema	166
4.1.91. io.k8s.api.core.v1.PodList schema	166
Schema	167
4.1.92. io.k8s.api.core.v1.PodOS schema	167
Schema	167
4.1.93. io.k8s.api.core.v1.PodReadinessGate schema	168
Schema	168
4.1.94. io.k8s.api.core.v1.PodResourceClaim schema	168
Schema	168
4.1.95. io.k8s.api.core.v1.PodSchedulingGate schema	169
Schema	169
4.1.96. io.k8s.api.core.v1.PodSecurityContext schema	169
Schema	169
4.1.97. io.k8s.api.core.v1.PodSpec schema	173
Schema	173
4.1.98. io.k8s.api.core.v1.PodTemplateList schema	183
Schema	184
4.1.99. io.k8s.api.core.v1.PodTemplateSpec schema	184
Schema	184
4.1.100. io.k8s.api.core.v1.PortworxVolumeSource schema	185
Schema	185
4.1.101. io.k8s.api.core.v1.PreferredSchedulingTerm schema	185
Schema	186
4.1.102. io.k8s.api.core.v1.Probe schema	186
Schema	186
4.1.103. io.k8s.api.core.v1.ProjectedVolumeSource schema	188
Schema	188
4.1.104. io.k8s.api.core.v1.QuobyteVolumeSource schema	188
Schema	189
4.1.105. io.k8s.api.core.v1.RBDPersistentVolumeSource schema	189
Schema	189
4.1.106. io.k8s.api.core.v1.RBDVolumeSource schema	191
Schema	191
4.1.107. io.k8s.api.core.v1.ReplicationControllerList schema	192
Schema	192
4.1.108. io.k8s.api.core.v1.ResourceClaim schema	193
Schema	193
4.1.109. io.k8s.api.core.v1.ResourceFieldSelector schema	193
Schema	194
4.1.110. io.k8s.api.core.v1.ResourceQuotaList schema	194
Schema	194
4.1.111. io.k8s.api.core.v1.ResourceRequirements schema	195
Schema	195
4.1.112. io.k8s.api.core.v1.ScaleIOPersistentVolumeSource schema	196
Schema	196
4.1.113. io.k8s.api.core.v1.ScaleIOVolumeSource schema	197
Schema	198
4.1.114. io.k8s.api.core.v1.SeccompProfile schema	199
Schema	199
4.1.115. io.k8s.api.core.v1.SecretEnvSource schema	200
Schema	200

4.1.116. io.k8s.api.core.v1.SecretKeySelector schema	200
Schema	201
4.1.117. io.k8s.api.core.v1.SecretList schema	201
Schema	201
4.1.118. io.k8s.api.core.v1.SecretProjection schema	202
Schema	202
4.1.119. io.k8s.api.core.v1.SecretReference schema	203
Schema	203
4.1.120. io.k8s.api.core.v1.SecretVolumeSource schema	203
Schema	204
4.1.121. io.k8s.api.core.v1.SecurityContext schema	205
Schema	205
4.1.122. io.k8s.api.core.v1.SELinuxOptions schema	208
Schema	208
4.1.123. io.k8s.api.core.v1.ServiceAccountList schema	208
Schema	208
4.1.124. io.k8s.api.core.v1.ServiceAccountTokenProjection schema	209
Schema	209
4.1.125. io.k8s.api.core.v1.ServiceList schema	210
Schema	210
4.1.126. io.k8s.api.core.v1.StorageOSPersistentVolumeSource schema	211
Schema	211
4.1.127. io.k8s.api.core.v1.StorageOSVolumeSource schema	212
Schema	212
4.1.128. io.k8s.api.core.v1.Sysctl schema	213
Schema	214
4.1.129. io.k8s.api.core.v1.TCPSocketAction schema	214
Schema	214
4.1.130. io.k8s.api.core.v1.Toleration schema	214
Schema	214
4.1.131. io.k8s.api.core.v1.TopologySelectorLabelRequirement schema	216
Schema	216
4.1.132. io.k8s.api.core.v1.TopologySelectorTerm schema	216
Schema	217
4.1.133. io.k8s.api.core.v1.TopologySpreadConstraint schema	217
Schema	217
4.1.134. io.k8s.api.core.v1.TypedLocalObjectReference schema	223
Schema	224
4.1.135. io.k8s.api.core.v1.TypedObjectReference schema	224
Schema	224
4.1.136. io.k8s.api.core.v1.Volume schema	225
Schema	225
4.1.137. io.k8s.api.core.v1.VolumeDevice schema	230
Schema	230
4.1.138. io.k8s.api.core.v1.VolumeMount schema	230
Schema	231
4.1.139. io.k8s.api.core.v1.VolumeNodeAffinity schema	232
Schema	232
4.1.140. io.k8s.api.core.v1.VolumeProjection schema	232
Schema	232
4.1.141. io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource schema	233
Schema	233
4.1.142. io.k8s.api.core.v1.WeightedPodAffinityTerm schema	233

Schema	234
4.1.143. io.k8s.api.core.v1.WindowsSecurityContextOptions schema	234
Schema	234
4.1.144. io.k8s.api.discovery.v1.EndpointSliceList schema	235
Schema	235
4.1.145. io.k8s.api.events.v1.EventList schema	236
Schema	236
4.1.146. io.k8s.api.flowcontrol.v1beta3.FlowSchemaList schema	237
Schema	237
4.1.147. io.k8s.api.flowcontrol.v1beta3.PriorityLevelConfigurationList schema	238
Schema	238
4.1.148. io.k8s.api.networking.v1.IngressClassList schema	239
Schema	239
4.1.149. io.k8s.api.networking.v1.IngressList schema	240
Schema	240
4.1.150. io.k8s.api.networking.v1.NetworkPolicyList schema	241
Schema	241
4.1.151. io.k8s.api.node.v1.RuntimeClassList schema	242
Schema	242
4.1.152. io.k8s.api.policy.v1.PodDisruptionBudgetList schema	243
Schema	243
4.1.153. io.k8s.api.rbac.v1.ClusterRoleBindingList schema	244
Schema	244
4.1.154. io.k8s.api.rbac.v1.ClusterRoleList schema	245
Schema	245
4.1.155. io.k8s.api.rbac.v1.RoleBindingList schema	246
Schema	246
4.1.156. io.k8s.api.rbac.v1.RoleList schema	247
Schema	247
4.1.157. io.k8s.api.scheduling.v1.PriorityClassList schema	248
Schema	248
4.1.158. io.k8s.api.storage.v1.CSIDriverList schema	249
Schema	249
4.1.159. io.k8s.api.storage.v1.CSINodeList schema	250
Schema	250
4.1.160. io.k8s.api.storage.v1.CSIStorageCapacityList schema	251
Schema	251
4.1.161. io.k8s.api.storage.v1.StorageClassList schema	252
Schema	252
4.1.162. io.k8s.api.storage.v1.VolumeAttachmentList schema	253
Schema	253
4.1.163. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.CustomResourceDefinitionList schema	254
Schema	254
4.1.164. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.ExternalDocumentation schema	255
Schema	255
4.1.165. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSON schema	256
4.1.166. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSONSchemaProps schema	256
Schema	256
4.1.167. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSONSchemaPropsOrArray schema	262
4.1.168. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSONSchemaPropsOrBool schema	262
4.1.169. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSONSchemaPropsOrStringArray schema	262
4.1.170. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.ValidationRule schema	263

Schema	263
4.1.171. io.k8s.apimachinery.pkg.api.resource.Quantity schema	266
4.1.172. io.k8s.apimachinery.pkg.apis.meta.v1.Condition schema	268
Schema	268
4.1.173. io.k8s.apimachinery.pkg.apis.meta.v1.DeleteOptions schema	269
Schema	269
4.1.174. io.k8s.apimachinery.pkg.apis.meta.v1.FieldsV1 schema	271
4.1.175. io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector schema	271
Schema	271
4.1.176. io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement schema	272
Schema	272
4.1.177. io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta schema	272
Schema	273
4.1.178. io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry schema	274
Schema	274
4.1.179. io.k8s.apimachinery.pkg.apis.meta.v1.MicroTime schema	275
4.1.180. io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta schema	275
Schema	276
4.1.181. io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference schema	282
Schema	282
4.1.182. io.k8s.apimachinery.pkg.apis.meta.v1.Patch schema	283
4.1.183. io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions schema	283
Schema	283
4.1.184. io.k8s.apimachinery.pkg.apis.meta.v1.Status schema	283
Schema	283
4.1.185. io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause schema	285
Schema	285
4.1.186. io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails schema	285
Schema	286
4.1.187. io.k8s.apimachinery.pkg.apis.meta.v1.Time schema	286
4.1.188. io.k8s.apimachinery.pkg.apis.meta.v1.WatchEvent schema	287
Schema	287
4.1.189. io.k8s.apimachinery.pkg.runtime.RawExtension schema	287
4.1.190. io.k8s.apimachinery.pkg.util.intstr.IntOrString schema	288
4.1.191. io.k8s.kube-agggregator.pkg.apis.apiregistration.v1.APIServiceList schema	288
Schema	288
4.1.192. io.k8s.migration.v1alpha1.StorageVersionMigrationList schema	289
Schema	289
4.1.193. io.k8s.storage.snapshot.v1.VolumeSnapshotClassList schema	290
Schema	290
4.1.194. io.k8s.storage.snapshot.v1.VolumeSnapshotContentList schema	291
Schema	291
4.1.195. io.k8s.storage.snapshot.v1.VolumeSnapshotList schema	292
Schema	292
4.1.196. io.openshift.internal.security.v1.RangeAllocationList schema	293
Schema	293
4.1.197. io.openshift.route.v1.RouteList schema	294
Schema	294
4.1.198. io.openshift.security.v1.SecurityContextConstraintsList schema	295
Schema	295
4.1.199. io.topolvm.v1.LogicalVolumeList schema	296
Schema	296

<b>CHAPTER 5. EXTENSION APIS</b> .....	<b>298</b>
5.1. API EXTENSIONS APIS	298
5.1.1. CustomResourceDefinition [apiextensions.k8s.io/v1]	298
5.2. CUSTOMRESOURCEDEFINITION [APIEXTENSIONS.K8S.IO/V1]	298
5.2.1. Specification	298
5.2.1.1. .spec	299
5.2.1.2. .spec.conversion	301
5.2.1.3. .spec.conversion.webhook	302
5.2.1.4. .spec.conversion.webhook.clientConfig	303
5.2.1.5. .spec.conversion.webhook.clientConfig.service	304
5.2.1.6. .spec.names	305
5.2.1.7. .spec.versions	306
5.2.1.8. .spec.versions[]	306
5.2.1.9. .spec.versions[].additionalPrinterColumns	308
5.2.1.10. .spec.versions[].additionalPrinterColumns[]	308
5.2.1.11. .spec.versions[].schema	309
5.2.1.12. .spec.versions[].subresources	310
5.2.1.13. .spec.versions[].subresources.scale	310
5.2.1.14. .spec.versions[].subresources.status	311
5.2.1.15. .status	312
5.2.1.16. .status.acceptedNames	312
5.2.1.17. .status.conditions	313
5.2.1.18. .status.conditions[]	314
5.2.2. API endpoints	314
5.2.2.1. /apis/apiextensions.k8s.io/v1/customresourcedefinitions	315
5.2.2.2. /apis/apiextensions.k8s.io/v1/watch/customresourcedefinitions	324
5.2.2.3. /apis/apiextensions.k8s.io/v1/customresourcedefinitions/{name}	327
5.2.2.4. /apis/apiextensions.k8s.io/v1/watch/customresourcedefinitions/{name}	331
5.2.2.5. /apis/apiextensions.k8s.io/v1/customresourcedefinitions/{name}/status	335
<b>CHAPTER 6. REGISTRATION APIS</b> .....	<b>339</b>
6.1. API REGISTRATION APIS	339
6.1.1. APIService [apiregistration.k8s.io/v1]	339
6.2. APISERVICE [APIREGISTRATION.K8S.IO/V1]	339
6.2.1. Specification	339
6.2.1.1. .spec	340
6.2.1.2. .spec.service	342
6.2.1.3. .status	343
6.2.1.4. .status.conditions	343
6.2.1.5. .status.conditions[]	343
6.2.2. API endpoints	344
6.2.2.1. /apis/apiregistration.k8s.io/v1/apiservices	345
6.2.2.2. /apis/apiregistration.k8s.io/v1/watch/apiservices	353
6.2.2.3. /apis/apiregistration.k8s.io/v1/apiservices/{name}	356
6.2.2.4. /apis/apiregistration.k8s.io/v1/watch/apiservices/{name}	360
6.2.2.5. /apis/apiregistration.k8s.io/v1/apiservices/{name}/status	364
<b>CHAPTER 7. APPS APIS</b> .....	<b>368</b>
7.1. APPS APIS	368
7.1.1. ControllerRevision [apps/v1]	368
7.1.2. DaemonSet [apps/v1]	368
7.1.3. Deployment [apps/v1]	368
7.1.4. ReplicaSet [apps/v1]	368

7.1.5. StatefulSet [apps/v1]	368
7.2. CONTROLLERREVISION [APPS/V1]	369
7.2.1. Specification	369
7.2.2. API endpoints	370
7.2.2.1. /apis/apps/v1/controllerrevisions	370
7.2.2.2. /apis/apps/v1/watch/controllerrevisions	374
7.2.2.3. /apis/apps/v1/namespaces/{namespace}/controllerrevisions	377
7.2.2.4. /apis/apps/v1/watch/namespaces/{namespace}/controllerrevisions	386
7.2.2.5. /apis/apps/v1/namespaces/{namespace}/controllerrevisions/{name}	390
7.2.2.6. /apis/apps/v1/watch/namespaces/{namespace}/controllerrevisions/{name}	394
7.3. DAEMONSET [APPS/V1]	398
7.3.1. Specification	398
7.3.1.1. .spec	399
7.3.1.2. .spec.updateStrategy	400
7.3.1.3. .spec.updateStrategy.rollingUpdate	400
7.3.1.4. .status	402
7.3.1.5. .status.conditions	404
7.3.1.6. .status.conditions[]	404
7.3.2. API endpoints	405
7.3.2.1. /apis/apps/v1/daemonsets	406
7.3.2.2. /apis/apps/v1/watch/daemonsets	409
7.3.2.3. /apis/apps/v1/namespaces/{namespace}/daemonsets	412
7.3.2.4. /apis/apps/v1/watch/namespaces/{namespace}/daemonsets	421
7.3.2.5. /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}	425
7.3.2.6. /apis/apps/v1/watch/namespaces/{namespace}/daemonsets/{name}	429
7.3.2.7. /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status	433
7.4. DEPLOYMENT [APPS/V1]	436
7.4.1. Specification	436
7.4.1.1. .spec	437
7.4.1.2. .spec.strategy	438
7.4.1.3. .spec.strategy.rollingUpdate	439
7.4.1.4. .status	440
7.4.1.5. .status.conditions	441
7.4.1.6. .status.conditions[]	441
7.4.2. API endpoints	442
7.4.2.1. /apis/apps/v1/deployments	443
7.4.2.2. /apis/apps/v1/watch/deployments	447
7.4.2.3. /apis/apps/v1/namespaces/{namespace}/deployments	450
7.4.2.4. /apis/apps/v1/watch/namespaces/{namespace}/deployments	459
7.4.2.5. /apis/apps/v1/namespaces/{namespace}/deployments/{name}	463
7.4.2.6. /apis/apps/v1/watch/namespaces/{namespace}/deployments/{name}	467
7.4.2.7. /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status	471
7.5. REPLICASET [APPS/V1]	474
7.5.1. Specification	474
7.5.1.1. .spec	475
7.5.1.2. .status	476
7.5.1.3. .status.conditions	477
7.5.1.4. .status.conditions[]	477
7.5.2. API endpoints	478
7.5.2.1. /apis/apps/v1/replicasets	479
7.5.2.2. /apis/apps/v1/watch/replicasets	482
7.5.2.3. /apis/apps/v1/namespaces/{namespace}/replicasets	485
7.5.2.4. /apis/apps/v1/watch/namespaces/{namespace}/replicasets	494

7.5.2.5. /apis/apps/v1/namespaces/{namespace}/replicasets/{name}	498
7.5.2.6. /apis/apps/v1/watch/namespaces/{namespace}/replicasets/{name}	502
7.5.2.7. /apis/apps/v1/namespaces/{namespace}/replicasets/{name}/status	506
7.6. STATEFULSET [APPS/V1]	509
7.6.1. Specification	509
7.6.1.1. .spec	510
7.6.1.2. .spec.ordinals	513
7.6.1.3. .spec.persistentVolumeClaimRetentionPolicy	514
7.6.1.4. .spec.updateStrategy	515
7.6.1.5. .spec.updateStrategy.rollingUpdate	515
7.6.1.6. .status	516
7.6.1.7. .status.conditions	518
7.6.1.8. .status.conditions[]	518
7.6.2. API endpoints	518
7.6.2.1. /apis/apps/v1/statefulsets	519
7.6.2.2. /apis/apps/v1/watch/statefulsets	523
7.6.2.3. /apis/apps/v1/namespaces/{namespace}/statefulsets	526
7.6.2.4. /apis/apps/v1/watch/namespaces/{namespace}/statefulsets	535
7.6.2.5. /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}	539
7.6.2.6. /apis/apps/v1/watch/namespaces/{namespace}/statefulsets/{name}	543
7.6.2.7. /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status	547
<b>CHAPTER 8. AUTHENTICATION APIS</b>	<b>551</b>
8.1. AUTHENTICATION APIS	551
8.1.1. TokenRequest [authentication.k8s.io/v1]	551
8.1.2. TokenReview [authentication.k8s.io/v1]	551
8.2. TOKENREQUEST [AUTHENTICATION.K8S.IO/V1]	551
8.2.1. Specification	551
8.2.1.1. .spec	552
8.2.1.2. .spec.boundObjectRef	553
8.2.1.3. .status	553
8.2.2. API endpoints	554
8.2.2.1. /api/v1/namespaces/{namespace}/serviceaccounts/{name}/token	554
8.3. TOKENREVIEW [AUTHENTICATION.K8S.IO/V1]	556
8.3.1. Specification	556
8.3.1.1. .spec	557
8.3.1.2. .status	557
8.3.1.3. .status.user	558
8.3.1.4. .status.user.extra	559
8.3.2. API endpoints	559
8.3.2.1. /apis/authentication.k8s.io/v1/tokenreviews	559
<b>CHAPTER 9. AUTHORIZATION APIS</b>	<b>561</b>
9.1. AUTHORIZATION APIS	561
9.1.1. LocalSubjectAccessReview [authorization.k8s.io/v1]	561
9.1.2. SelfSubjectAccessReview [authorization.k8s.io/v1]	561
9.1.3. SelfSubjectRulesReview [authorization.k8s.io/v1]	561
9.1.4. SubjectAccessReview [authorization.k8s.io/v1]	561
9.2. LOCALSUBJECTACCESSREVIEW [AUTHORIZATION.K8S.IO/V1]	561
9.2.1. Specification	562
9.2.1.1. .spec	562
9.2.1.2. .spec.extra	563
9.2.1.3. .spec.nonResourceAttributes	563

9.2.1.4. .spec.resourceAttributes	564
9.2.1.5. .status	565
9.2.2. API endpoints	566
9.2.2.1. /apis/authorization.k8s.io/v1/namespaces/{namespace}/localsubjectaccessreviews	566
9.3. SELFSUBJECTACCESSREVIEW [AUTHORIZATION.K8S.IO/V1]	567
9.3.1. Specification	567
9.3.1.1. .spec	568
9.3.1.2. .spec.nonResourceAttributes	569
9.3.1.3. .spec.resourceAttributes	569
9.3.1.4. .status	570
9.3.2. API endpoints	571
9.3.2.1. /apis/authorization.k8s.io/v1/selfsubjectaccessreviews	571
9.4. SELFSUBJECTRULESREVIEW [AUTHORIZATION.K8S.IO/V1]	573
9.4.1. Specification	573
9.4.1.1. .spec	574
9.4.1.2. .status	574
9.4.1.3. .status.nonResourceRules	576
9.4.1.4. .status.nonResourceRules[]	576
9.4.1.5. .status.resourceRules	576
9.4.1.6. .status.resourceRules[]	576
9.4.2. API endpoints	577
9.4.2.1. /apis/authorization.k8s.io/v1/selfsubjectrulesreviews	577
9.5. SUBJECTACCESSREVIEW [AUTHORIZATION.K8S.IO/V1]	579
9.5.1. Specification	579
9.5.1.1. .spec	580
9.5.1.2. .spec.extra	581
9.5.1.3. .spec.nonResourceAttributes	581
9.5.1.4. .spec.resourceAttributes	581
9.5.1.5. .status	582
9.5.2. API endpoints	583
9.5.2.1. /apis/authorization.k8s.io/v1/subjectaccessreviews	583
<b>CHAPTER 10. AUTOSCALING APIS .....</b>	<b>586</b>
10.1. AUTOSCALING APIS	586
10.1.1. Scale [autoscaling/v1]	586
10.1.2. HorizontalPodAutoscaler [autoscaling/v2]	586
10.2. HORIZONTALPODAUTOSCALER [AUTOSCALING/V2]	586
10.2.1. Specification	586
10.2.1.1. .spec	587
10.2.1.2. .spec.behavior	588
10.2.1.3. .spec.behavior.scaleDown	589
10.2.1.4. .spec.behavior.scaleDown.policies	590
10.2.1.5. .spec.behavior.scaleDown.policies[]	590
10.2.1.6. .spec.behavior.scaleUp	591
10.2.1.7. .spec.behavior.scaleUp.policies	592
10.2.1.8. .spec.behavior.scaleUp.policies[]	592
10.2.1.9. .spec.metrics	593
10.2.1.10. .spec.metrics[]	593
10.2.1.11. .spec.metrics[].containerResource	595
10.2.1.12. .spec.metrics[].containerResource.target	596
10.2.1.13. .spec.metrics[].external	597
10.2.1.14. .spec.metrics[].external.metric	597
10.2.1.15. .spec.metrics[].external.target	598



10.2.1.16. .spec.metrics[].object	598
10.2.1.17. .spec.metrics[].object.describedObject	599
10.2.1.18. .spec.metrics[].object.metric	599
10.2.1.19. .spec.metrics[].object.target	600
10.2.1.20. .spec.metrics[].pods	601
10.2.1.21. .spec.metrics[].pods.metric	601
10.2.1.22. .spec.metrics[].pods.target	602
10.2.1.23. .spec.metrics[].resource	603
10.2.1.24. .spec.metrics[].resource.target	603
10.2.1.25. .spec.scaleTargetRef	604
10.2.1.26. .status	604
10.2.1.27. .status.conditions	605
10.2.1.28. .status.conditions[]	606
10.2.1.29. .status.currentMetrics	606
10.2.1.30. .status.currentMetrics[]	606
10.2.1.31. .status.currentMetrics[].containerResource	608
10.2.1.32. .status.currentMetrics[].containerResource.current	608
10.2.1.33. .status.currentMetrics[].external	609
10.2.1.34. .status.currentMetrics[].external.current	609
10.2.1.35. .status.currentMetrics[].external.metric	610
10.2.1.36. .status.currentMetrics[].object	610
10.2.1.37. .status.currentMetrics[].object.current	611
10.2.1.38. .status.currentMetrics[].object.describedObject	611
10.2.1.39. .status.currentMetrics[].object.metric	612
10.2.1.40. .status.currentMetrics[].pods	613
10.2.1.41. .status.currentMetrics[].pods.current	613
10.2.1.42. .status.currentMetrics[].pods.metric	614
10.2.1.43. .status.currentMetrics[].resource	614
10.2.1.44. .status.currentMetrics[].resource.current	615
10.2.2. API endpoints	615
10.2.2.1. /apis/autoscaling/v2/horizontalpodautoscalers	616
10.2.2.2. /apis/autoscaling/v2/watch/horizontalpodautoscalers	620
10.2.2.3. /apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers	623
10.2.2.4. /apis/autoscaling/v2/watch/namespaces/{namespace}/horizontalpodautoscalers	632
10.2.2.5. /apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers/{name}	636
10.2.2.6. /apis/autoscaling/v2/watch/namespaces/{namespace}/horizontalpodautoscalers/{name}	640
10.2.2.7. /apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers/{name}/status	644
10.3. SCALE [AUTOSCALING/V1]	647
10.3.1. Specification	647
10.3.1.1. .spec	648
10.3.1.2. .status	648
10.3.2. API endpoints	649
10.3.2.1. /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale	649
10.3.2.2. /apis/apps/v1/namespaces/{namespace}/replicasets/{name}/scale	653
10.3.2.3. /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/scale	656
10.3.2.4. /api/v1/namespaces/{namespace}/replicationcontrollers/{name}/scale	660
<b>CHAPTER 11. BATCH APIS</b>	<b>664</b>
11.1. BATCH APIS	664
11.1.1. CronJob [batch/v1]	664
11.1.2. Job [batch/v1]	664
11.2. CRONJOB [BATCH/V1]	664
11.2.1. Specification	664

11.2.1.1. .spec	665
11.2.1.2. .spec.jobTemplate	667
11.2.1.3. .spec.jobTemplate.spec	668
11.2.1.4. .spec.jobTemplate.spec.podFailurePolicy	671
11.2.1.5. .spec.jobTemplate.spec.podFailurePolicy.rules	672
11.2.1.6. .spec.jobTemplate.spec.podFailurePolicy.rules[]	672
11.2.1.7. .spec.jobTemplate.spec.podFailurePolicy.rules[].onExitCodes	674
11.2.1.8. .spec.jobTemplate.spec.podFailurePolicy.rules[].onPodConditions	675
11.2.1.9. .spec.jobTemplate.spec.podFailurePolicy.rules[].onPodConditions[]	676
11.2.1.10. .status	676
11.2.2. API endpoints	677
11.2.2.1. /apis/batch/v1/cronjobs	678
11.2.2.2. /apis/batch/v1/watch/cronjobs	681
11.2.2.3. /apis/batch/v1/namespaces/{namespace}/cronjobs	684
11.2.2.4. /apis/batch/v1/watch/namespaces/{namespace}/cronjobs	693
11.2.2.5. /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}	697
11.2.2.6. /apis/batch/v1/watch/namespaces/{namespace}/cronjobs/{name}	701
11.2.2.7. /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status	705
11.3. JOB [BATCH/V1]	708
11.3.1. Specification	708
11.3.1.1. .spec	709
11.3.1.2. .spec.podFailurePolicy	713
11.3.1.3. .spec.podFailurePolicy.rules	714
11.3.1.4. .spec.podFailurePolicy.rules[]	714
11.3.1.5. .spec.podFailurePolicy.rules[].onExitCodes	716
11.3.1.6. .spec.podFailurePolicy.rules[].onPodConditions	717
11.3.1.7. .spec.podFailurePolicy.rules[].onPodConditions[]	718
11.3.1.8. .status	718
11.3.1.9. .status.conditions	720
11.3.1.10. .status.conditions[]	720
11.3.1.11. .status.uncountedTerminatedPods	721
11.3.2. API endpoints	721
11.3.2.1. /apis/batch/v1/jobs	722
11.3.2.2. /apis/batch/v1/watch/jobs	726
11.3.2.3. /apis/batch/v1/namespaces/{namespace}/jobs	729
11.3.2.4. /apis/batch/v1/watch/namespaces/{namespace}/jobs	738
11.3.2.5. /apis/batch/v1/namespaces/{namespace}/jobs/{name}	742
11.3.2.6. /apis/batch/v1/watch/namespaces/{namespace}/jobs/{name}	746
11.3.2.7. /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status	750
<b>CHAPTER 12. CERTIFICATES APIS</b>	<b>754</b>
12.1. CERTIFICATES APIS	754
12.1.1. CertificateSigningRequest [certificates.k8s.io/v1]	754
12.2. CERTIFICATESIGNINGREQUEST [CERTIFICATES.K8S.IO/V1]	754
12.2.1. Specification	754
12.2.1.1. .spec	755
12.2.1.2. .spec.extra	759
12.2.1.3. .status	759
12.2.1.4. .status.conditions	761
12.2.1.5. .status.conditions[]	761
12.2.2. API endpoints	762
12.2.2.1. /apis/certificates.k8s.io/v1/certificatesigningrequests	763
12.2.2.2. /apis/certificates.k8s.io/v1/watch/certificatesigningrequests	772

12.2.2.3. /apis/certificates.k8s.io/v1/certificatesigningrequests/{name}	775
12.2.2.4. /apis/certificates.k8s.io/v1/watch/certificatesigningrequests/{name}	779
12.2.2.5. /apis/certificates.k8s.io/v1/certificatesigningrequests/{name}/status	783
12.2.2.6. /apis/certificates.k8s.io/v1/certificatesigningrequests/{name}/approval	786
<b>CHAPTER 13. COORDINATION APIS</b> .....	<b>790</b>
13.1. COORDINATION APIS	790
13.1.1. Lease [coordination.k8s.io/v1]	790
13.2. LEASE [COORDINATION.K8S.IO/V1]	790
13.2.1. Specification	790
13.2.1.1. .spec	791
13.2.2. API endpoints	791
13.2.2.1. /apis/coordination.k8s.io/v1/leases	792
13.2.2.2. /apis/coordination.k8s.io/v1/watch/leases	796
13.2.2.3. /apis/coordination.k8s.io/v1/namespaces/{namespace}/leases	799
13.2.2.4. /apis/coordination.k8s.io/v1/watch/namespaces/{namespace}/leases	808
13.2.2.5. /apis/coordination.k8s.io/v1/namespaces/{namespace}/leases/{name}	812
13.2.2.6. /apis/coordination.k8s.io/v1/watch/namespaces/{namespace}/leases/{name}	816
<b>CHAPTER 14. CORE APIS</b> .....	<b>821</b>
14.1. CORE APIS	821
14.1.1. Binding [v1]	821
14.1.2. ComponentStatus [v1]	821
14.1.3. ConfigMap [v1]	821
14.1.4. Endpoints [v1]	821
14.1.5. Event [v1]	821
14.1.6. LimitRange [v1]	822
14.1.7. Namespace [v1]	822
14.1.8. Node [v1]	822
14.1.9. PersistentVolume [v1]	822
14.1.10. PersistentVolumeClaim [v1]	822
14.1.11. Pod [v1]	822
14.1.12. PodTemplate [v1]	823
14.1.13. ReplicationController [v1]	823
14.1.14. ResourceQuota [v1]	823
14.1.15. Secret [v1]	823
14.1.16. Service [v1]	823
14.1.17. ServiceAccount [v1]	823
14.2. BINDING [V1]	824
14.2.1. Specification	824
14.2.1.1. .target	825
14.2.2. API endpoints	826
14.2.2.1. /api/v1/namespaces/{namespace}/bindings	826
14.2.2.2. /api/v1/namespaces/{namespace}/pods/{name}/binding	828
14.3. COMPONENTSTATUS [V1]	829
14.3.1. Specification	830
14.3.1.1. .conditions	830
14.3.1.2. .conditions[]	831
14.3.2. API endpoints	831
14.3.2.1. /api/v1/componentstatuses	831
14.3.2.2. /api/v1/componentstatuses/{name}	835
14.4. CONFIGMAP [V1]	835
14.4.1. Specification	835

14.4.2. API endpoints	837
14.4.2.1. /api/v1/configmaps	838
14.4.2.2. /api/v1/watch/configmaps	841
14.4.2.3. /api/v1/namespaces/{namespace}/configmaps	844
14.4.2.4. /api/v1/watch/namespaces/{namespace}/configmaps	853
14.4.2.5. /api/v1/namespaces/{namespace}/configmaps/{name}	857
14.4.2.6. /api/v1/watch/namespaces/{namespace}/configmaps/{name}	861
14.5. ENDPOINTS [V1]	865
14.5.1. Specification	865
14.5.1.1. .subsets	867
14.5.1.2. .subsets[]	867
14.5.1.3. .subsets[].addresses	868
14.5.1.4. .subsets[].addresses[]	868
14.5.1.5. .subsets[].addresses[].targetRef	869
14.5.1.6. .subsets[].notReadyAddresses	871
14.5.1.7. .subsets[].notReadyAddresses[]	871
14.5.1.8. .subsets[].notReadyAddresses[].targetRef	871
14.5.1.9. .subsets[].ports	873
14.5.1.10. .subsets[].ports[]	873
14.5.2. API endpoints	874
14.5.2.1. /api/v1/endpoints	875
14.5.2.2. /api/v1/watch/endpoints	879
14.5.2.3. /api/v1/namespaces/{namespace}/endpoints	882
14.5.2.4. /api/v1/watch/namespaces/{namespace}/endpoints	891
14.5.2.5. /api/v1/namespaces/{namespace}/endpoints/{name}	895
14.5.2.6. /api/v1/watch/namespaces/{namespace}/endpoints/{name}	899
14.6. EVENT [V1]	903
14.6.1. Specification	903
14.6.1.1. .involvedObject	905
14.6.1.2. .related	906
14.6.1.3. .series	908
14.6.1.4. .source	908
14.6.2. API endpoints	908
14.6.2.1. /api/v1/events	909
14.6.2.2. /api/v1/watch/events	913
14.6.2.3. /api/v1/namespaces/{namespace}/events	916
14.6.2.4. /api/v1/watch/namespaces/{namespace}/events	925
14.6.2.5. /api/v1/namespaces/{namespace}/events/{name}	929
14.6.2.6. /api/v1/watch/namespaces/{namespace}/events/{name}	933
14.7. LIMITRANGE [V1]	937
14.7.1. Specification	937
14.7.1.1. .spec	938
14.7.1.2. .spec.limits	938
14.7.1.3. .spec.limits[]	938
14.7.2. API endpoints	939
14.7.2.1. /api/v1/limitranges	940
14.7.2.2. /api/v1/watch/limitranges	944
14.7.2.3. /api/v1/namespaces/{namespace}/limitranges	947
14.7.2.4. /api/v1/watch/namespaces/{namespace}/limitranges	956
14.7.2.5. /api/v1/namespaces/{namespace}/limitranges/{name}	960
14.7.2.6. /api/v1/watch/namespaces/{namespace}/limitranges/{name}	964
14.8. NAMESPACE [V1]	968
14.8.1. Specification	968

14.8.1.1. .spec	969
14.8.1.2. .status	969
14.8.1.3. .status.conditions	970
14.8.1.4. .status.conditions[]	970
14.8.2. API endpoints	971
14.8.2.1. /api/v1/namespaces	971
14.8.2.2. /api/v1/watch/namespaces	976
14.8.2.3. /api/v1/namespaces/{name}	979
14.8.2.4. /api/v1/watch/namespaces/{name}	983
14.8.2.5. /api/v1/namespaces/{name}/status	987
14.8.2.6. /api/v1/namespaces/{name}/finalize	990
14.9. NODE [V1]	991
14.9.1. Specification	992
14.9.1.1. .spec	992
14.9.1.2. .spec.configSource	993
14.9.1.3. .spec.configSource.configMap	994
14.9.1.4. .spec.taints	995
14.9.1.5. .spec.taints[]	995
14.9.1.6. .status	996
14.9.1.7. .status.addresses	998
14.9.1.8. .status.addresses[]	999
14.9.1.9. .status.conditions	999
14.9.1.10. .status.conditions[]	999
14.9.1.11. .status.config	1000
14.9.1.12. .status.config.active	1001
14.9.1.13. .status.config.active.configMap	1002
14.9.1.14. .status.config.assigned	1003
14.9.1.15. .status.config.assigned.configMap	1003
14.9.1.16. .status.config.lastKnownGood	1004
14.9.1.17. .status.config.lastKnownGood.configMap	1004
14.9.1.18. .status.daemonEndpoints	1005
14.9.1.19. .status.daemonEndpoints.kubeletEndpoint	1006
14.9.1.20. .status.images	1006
14.9.1.21. .status.images[]	1006
14.9.1.22. .status.nodeInfo	1007
14.9.1.23. .status.volumesAttached	1008
14.9.1.24. .status.volumesAttached[]	1008
14.9.2. API endpoints	1009
14.9.2.1. /api/v1/nodes	1009
14.9.2.2. /api/v1/watch/nodes	1018
14.9.2.3. /api/v1/nodes/{name}	1021
14.9.2.4. /api/v1/watch/nodes/{name}	1025
14.9.2.5. /api/v1/nodes/{name}/status	1029
14.10. PERSISTENTVOLUME [V1]	1032
14.10.1. Specification	1032
14.10.1.1. .spec	1033
14.10.1.2. .spec.awsElasticBlockStore	1038
14.10.1.3. .spec.azureDisk	1039
14.10.1.4. .spec.azureFile	1040
14.10.1.5. .spec.cephfs	1041
14.10.1.6. .spec.cephfs.secretRef	1042
14.10.1.7. .spec.cinder	1042
14.10.1.8. .spec.cinder.secretRef	1043

14.10.1.9. .spec.claimRef	1043
14.10.1.10. .spec.csi	1045
14.10.1.11. .spec.csi.controllerExpandSecretRef	1046
14.10.1.12. .spec.csi.controllerPublishSecretRef	1047
14.10.1.13. .spec.csi.nodeExpandSecretRef	1047
14.10.1.14. .spec.csi.nodePublishSecretRef	1047
14.10.1.15. .spec.csi.nodeStageSecretRef	1048
14.10.1.16. .spec.fc	1048
14.10.1.17. .spec.flexVolume	1049
14.10.1.18. .spec.flexVolume.secretRef	1050
14.10.1.19. .spec.flocker	1050
14.10.1.20. .spec.gcePersistentDisk	1051
14.10.1.21. .spec.glusterfs	1052
14.10.1.22. .spec.hostPath	1053
14.10.1.23. .spec.iscsi	1054
14.10.1.24. .spec.iscsi.secretRef	1056
14.10.1.25. .spec.local	1056
14.10.1.26. .spec.nfs	1057
14.10.1.27. .spec.nodeAffinity	1057
14.10.1.28. .spec.nodeAffinity.required	1058
14.10.1.29. .spec.nodeAffinity.required.nodeSelectorTerms	1058
14.10.1.30. .spec.nodeAffinity.required.nodeSelectorTerms[]	1058
14.10.1.31. .spec.nodeAffinity.required.nodeSelectorTerms[].matchExpressions	1059
14.10.1.32. .spec.nodeAffinity.required.nodeSelectorTerms[].matchExpressions[]	1059
14.10.1.33. .spec.nodeAffinity.required.nodeSelectorTerms[].matchFields	1060
14.10.1.34. .spec.nodeAffinity.required.nodeSelectorTerms[].matchFields[]	1060
14.10.1.35. .spec.photonPersistentDisk	1061
14.10.1.36. .spec.portworxVolume	1061
14.10.1.37. .spec.quobyte	1062
14.10.1.38. .spec.rbd	1063
14.10.1.39. .spec.rbd.secretRef	1064
14.10.1.40. .spec.scaleIO	1065
14.10.1.41. .spec.scaleIO.secretRef	1066
14.10.1.42. .spec.storageos	1066
14.10.1.43. .spec.storageos.secretRef	1067
14.10.1.44. .spec.vsphereVolume	1069
14.10.1.45. .status	1069
14.10.2. API endpoints	1070
14.10.2.1. /api/v1/persistentvolumes	1071
14.10.2.2. /api/v1/watch/persistentvolumes	1080
14.10.2.3. /api/v1/persistentvolumes/{name}	1083
14.10.2.4. /api/v1/watch/persistentvolumes/{name}	1087
14.10.2.5. /api/v1/persistentvolumes/{name}/status	1091
14.11. PERSISTENTVOLUMECLAIM [V1]	1094
14.11.1. Specification	1094
14.11.1.1. .spec	1095
14.11.1.2. .spec.dataSource	1097
14.11.1.3. .spec.dataSourceRef	1098
14.11.1.4. .spec.resources	1099
14.11.1.5. .spec.resources.claims	1100
14.11.1.6. .spec.resources.claims[]	1100
14.11.1.7. .status	1101
14.11.1.8. .status.conditions	1103

14.11.1.9. .status.conditions[]	1104
14.11.2. API endpoints	1104
14.11.2.1. /api/v1/persistentvolumeclaims	1105
14.11.2.2. /api/v1/watch/persistentvolumeclaims	1109
14.11.2.3. /api/v1/namespaces/{namespace}/persistentvolumeclaims	1112
14.11.2.4. /api/v1/watch/namespaces/{namespace}/persistentvolumeclaims	1121
14.11.2.5. /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}	1125
14.11.2.6. /api/v1/watch/namespaces/{namespace}/persistentvolumeclaims/{name}	1129
14.11.2.7. /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}/status	1133
14.12. POD [V1]	1136
14.12.1. Specification	1136
14.12.1.1. .spec	1137
14.12.1.2. .spec.affinity	1149
14.12.1.3. .spec.affinity.nodeAffinity	1149
14.12.1.4. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution	1150
14.12.1.5. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1150
14.12.1.6. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference	1151
14.12.1.7. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchEx	1151
14.12.1.8. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchEx	1152
14.12.1.9. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFi	1152
14.12.1.10. .spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference.matchFi	1152
14.12.1.11. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution	1153
14.12.1.12. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms	1154
14.12.1.13. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]	1154
14.12.1.14. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].r	1154
14.12.1.15. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].r	1155
14.12.1.16. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].r	1155
14.12.1.17. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[].r	1156
14.12.1.18. .spec.affinity.podAffinity	1156
14.12.1.19. .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution	1158
14.12.1.20. .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1158
14.12.1.21. .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm	1159
14.12.1.22. .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution	1160
14.12.1.23. .spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]	1160
14.12.1.24. .spec.affinity.podAntiAffinity	1161
14.12.1.25. .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution	1163
14.12.1.26. .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1163
14.12.1.27. .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinityTerm	1164
14.12.1.28. .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution	1165
14.12.1.29. .spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]	1165
14.12.1.30. .spec.containers	1166
14.12.1.31. .spec.containers[]	1166

14.12.1.32. .spec.containers[].env	1173
14.12.1.33. .spec.containers[].env[]	1173
14.12.1.34. .spec.containers[].env[].valueFrom	1174
14.12.1.35. .spec.containers[].env[].valueFrom.configMapKeyRef	1174
14.12.1.36. .spec.containers[].env[].valueFrom.fieldRef	1175
14.12.1.37. .spec.containers[].env[].valueFrom.resourceFieldRef	1175
14.12.1.38. .spec.containers[].env[].valueFrom.secretKeyRef	1176
14.12.1.39. .spec.containers[].envFrom	1176
14.12.1.40. .spec.containers[].envFrom[]	1176
14.12.1.41. .spec.containers[].envFrom[].configMapRef	1177
14.12.1.42. .spec.containers[].envFrom[].secretRef	1177
14.12.1.43. .spec.containers[].lifecycle	1178
14.12.1.44. .spec.containers[].lifecycle.postStart	1178
14.12.1.45. .spec.containers[].lifecycle.postStart.exec	1179
14.12.1.46. .spec.containers[].lifecycle.postStart.httpGet	1179
14.12.1.47. .spec.containers[].lifecycle.postStart.httpGet.httpHeaders	1180
14.12.1.48. .spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]	1180
14.12.1.49. .spec.containers[].lifecycle.postStart.tcpSocket	1181
14.12.1.50. .spec.containers[].lifecycle.preStop	1181
14.12.1.51. .spec.containers[].lifecycle.preStop.exec	1182
14.12.1.52. .spec.containers[].lifecycle.preStop.httpGet	1182
14.12.1.53. .spec.containers[].lifecycle.preStop.httpGet.httpHeaders	1183
14.12.1.54. .spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]	1183
14.12.1.55. .spec.containers[].lifecycle.preStop.tcpSocket	1184
14.12.1.56. .spec.containers[].livenessProbe	1184
14.12.1.57. .spec.containers[].livenessProbe.exec	1186
14.12.1.58. .spec.containers[].livenessProbe.grpc	1187
14.12.1.59. .spec.containers[].livenessProbe.httpGet	1187
14.12.1.60. .spec.containers[].livenessProbe.httpGet.httpHeaders	1188
14.12.1.61. .spec.containers[].livenessProbe.httpGet.httpHeaders[]	1188
14.12.1.62. .spec.containers[].livenessProbe.tcpSocket	1189
14.12.1.63. .spec.containers[].ports	1189
14.12.1.64. .spec.containers[].ports[]	1189
14.12.1.65. .spec.containers[].readinessProbe	1190
14.12.1.66. .spec.containers[].readinessProbe.exec	1192
14.12.1.67. .spec.containers[].readinessProbe.grpc	1193
14.12.1.68. .spec.containers[].readinessProbe.httpGet	1193
14.12.1.69. .spec.containers[].readinessProbe.httpGet.httpHeaders	1194
14.12.1.70. .spec.containers[].readinessProbe.httpGet.httpHeaders[]	1194
14.12.1.71. .spec.containers[].readinessProbe.tcpSocket	1195
14.12.1.72. .spec.containers[].resizePolicy	1195
14.12.1.73. .spec.containers[].resizePolicy[]	1195
14.12.1.74. .spec.containers[].resources	1196
14.12.1.75. .spec.containers[].resources.claims	1197
14.12.1.76. .spec.containers[].resources.claims[]	1197
14.12.1.77. .spec.containers[].securityContext	1197
14.12.1.78. .spec.containers[].securityContext.capabilities	1200
14.12.1.79. .spec.containers[].securityContext.seLinuxOptions	1201
14.12.1.80. .spec.containers[].securityContext.seccompProfile	1201
14.12.1.81. .spec.containers[].securityContext.windowsOptions	1202
14.12.1.82. .spec.containers[].startupProbe	1203
14.12.1.83. .spec.containers[].startupProbe.exec	1205
14.12.1.84. .spec.containers[].startupProbe.grpc	1206



14.12.1.85. .spec.containers[].startupProbe.httpGet	1206
14.12.1.86. .spec.containers[].startupProbe.httpGet.httpHeaders	1207
14.12.1.87. .spec.containers[].startupProbe.httpGet.httpHeaders[]	1207
14.12.1.88. .spec.containers[].startupProbe.tcpSocket	1208
14.12.1.89. .spec.containers[].volumeDevices	1208
14.12.1.90. .spec.containers[].volumeDevices[]	1208
14.12.1.91. .spec.containers[].volumeMounts	1209
14.12.1.92. .spec.containers[].volumeMounts[]	1209
14.12.1.93. .spec.dnsConfig	1211
14.12.1.94. .spec.dnsConfig.options	1211
14.12.1.95. .spec.dnsConfig.options[]	1212
14.12.1.96. .spec.ephemeralContainers	1212
14.12.1.97. .spec.ephemeralContainers[]	1212
14.12.1.98. .spec.ephemeralContainers[].env	1219
14.12.1.99. .spec.ephemeralContainers[].env[]	1219
14.12.1.100. .spec.ephemeralContainers[].env[].valueFrom	1220
14.12.1.101. .spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef	1220
14.12.1.102. .spec.ephemeralContainers[].env[].valueFrom.fieldRef	1221
14.12.1.103. .spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef	1221
14.12.1.104. .spec.ephemeralContainers[].env[].valueFrom.secretKeyRef	1222
14.12.1.105. .spec.ephemeralContainers[].envFrom	1222
14.12.1.106. .spec.ephemeralContainers[].envFrom[]	1222
14.12.1.107. .spec.ephemeralContainers[].envFrom[].configMapRef	1223
14.12.1.108. .spec.ephemeralContainers[].envFrom[].secretRef	1223
14.12.1.109. .spec.ephemeralContainers[].lifecycle	1224
14.12.1.110. .spec.ephemeralContainers[].lifecycle.postStart	1224
14.12.1.111. .spec.ephemeralContainers[].lifecycle.postStart.exec	1225
14.12.1.112. .spec.ephemeralContainers[].lifecycle.postStart.httpGet	1225
14.12.1.113. .spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders	1226
14.12.1.114. .spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]	1226
14.12.1.115. .spec.ephemeralContainers[].lifecycle.postStart.tcpSocket	1227
14.12.1.116. .spec.ephemeralContainers[].lifecycle.preStop	1227
14.12.1.117. .spec.ephemeralContainers[].lifecycle.preStop.exec	1228
14.12.1.118. .spec.ephemeralContainers[].lifecycle.preStop.httpGet	1228
14.12.1.119. .spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders	1229
14.12.1.120. .spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]	1229
14.12.1.121. .spec.ephemeralContainers[].lifecycle.preStop.tcpSocket	1230
14.12.1.122. .spec.ephemeralContainers[].livenessProbe	1230
14.12.1.123. .spec.ephemeralContainers[].livenessProbe.exec	1232
14.12.1.124. .spec.ephemeralContainers[].livenessProbe.grpc	1233
14.12.1.125. .spec.ephemeralContainers[].livenessProbe.httpGet	1233
14.12.1.126. .spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders	1234
14.12.1.127. .spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]	1234
14.12.1.128. .spec.ephemeralContainers[].livenessProbe.tcpSocket	1235
14.12.1.129. .spec.ephemeralContainers[].ports	1235
14.12.1.130. .spec.ephemeralContainers[].ports[]	1235
14.12.1.131. .spec.ephemeralContainers[].readinessProbe	1236
14.12.1.132. .spec.ephemeralContainers[].readinessProbe.exec	1238
14.12.1.133. .spec.ephemeralContainers[].readinessProbe.grpc	1239
14.12.1.134. .spec.ephemeralContainers[].readinessProbe.httpGet	1239
14.12.1.135. .spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders	1240
14.12.1.136. .spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]	1240
14.12.1.137. .spec.ephemeralContainers[].readinessProbe.tcpSocket	1241

14.12.1.138. .spec.ephemeralContainers[].resizePolicy	1241
14.12.1.139. .spec.ephemeralContainers[].resizePolicy[]	1241
14.12.1.140. .spec.ephemeralContainers[].resources	1242
14.12.1.141. .spec.ephemeralContainers[].resources.claims	1243
14.12.1.142. .spec.ephemeralContainers[].resources.claims[]	1243
14.12.1.143. .spec.ephemeralContainers[].securityContext	1243
14.12.1.144. .spec.ephemeralContainers[].securityContext.capabilities	1246
14.12.1.145. .spec.ephemeralContainers[].securityContext.seLinuxOptions	1247
14.12.1.146. .spec.ephemeralContainers[].securityContext.seccompProfile	1247
14.12.1.147. .spec.ephemeralContainers[].securityContext.windowsOptions	1248
14.12.1.148. .spec.ephemeralContainers[].startupProbe	1249
14.12.1.149. .spec.ephemeralContainers[].startupProbe.exec	1251
14.12.1.150. .spec.ephemeralContainers[].startupProbe.grpc	1252
14.12.1.151. .spec.ephemeralContainers[].startupProbe.httpGet	1252
14.12.1.152. .spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders	1253
14.12.1.153. .spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]	1253
14.12.1.154. .spec.ephemeralContainers[].startupProbe.tcpSocket	1254
14.12.1.155. .spec.ephemeralContainers[].volumeDevices	1254
14.12.1.156. .spec.ephemeralContainers[].volumeDevices[]	1254
14.12.1.157. .spec.ephemeralContainers[].volumeMounts	1255
14.12.1.158. .spec.ephemeralContainers[].volumeMounts[]	1255
14.12.1.159. .spec.hostAliases	1257
14.12.1.160. .spec.hostAliases[]	1257
14.12.1.161. .spec.imagePullSecrets	1257
14.12.1.162. .spec.imagePullSecrets[]	1257
14.12.1.163. .spec.initContainers	1258
14.12.1.164. .spec.initContainers[]	1258
14.12.1.165. .spec.initContainers[].env	1265
14.12.1.166. .spec.initContainers[].env[]	1265
14.12.1.167. .spec.initContainers[].env[].valueFrom	1266
14.12.1.168. .spec.initContainers[].env[].valueFrom.configMapKeyRef	1266
14.12.1.169. .spec.initContainers[].env[].valueFrom.fieldRef	1267
14.12.1.170. .spec.initContainers[].env[].valueFrom.resourceFieldRef	1267
14.12.1.171. .spec.initContainers[].env[].valueFrom.secretKeyRef	1268
14.12.1.172. .spec.initContainers[].envFrom	1268
14.12.1.173. .spec.initContainers[].envFrom[]	1268
14.12.1.174. .spec.initContainers[].envFrom[].configMapRef	1269
14.12.1.175. .spec.initContainers[].envFrom[].secretRef	1270
14.12.1.176. .spec.initContainers[].lifecycle	1270
14.12.1.177. .spec.initContainers[].lifecycle.postStart	1270
14.12.1.178. .spec.initContainers[].lifecycle.postStart.exec	1271
14.12.1.179. .spec.initContainers[].lifecycle.postStart.httpGet	1271
14.12.1.180. .spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders	1272
14.12.1.181. .spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]	1272
14.12.1.182. .spec.initContainers[].lifecycle.postStart.tcpSocket	1273
14.12.1.183. .spec.initContainers[].lifecycle.preStop	1273
14.12.1.184. .spec.initContainers[].lifecycle.preStop.exec	1274
14.12.1.185. .spec.initContainers[].lifecycle.preStop.httpGet	1274
14.12.1.186. .spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders	1275
14.12.1.187. .spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]	1275
14.12.1.188. .spec.initContainers[].lifecycle.preStop.tcpSocket	1276
14.12.1.189. .spec.initContainers[].livenessProbe	1276
14.12.1.190. .spec.initContainers[].livenessProbe.exec	1278

14.12.1.191. .spec.initContainers[].livenessProbe.grpc	1279
14.12.1.192. .spec.initContainers[].livenessProbe.httpGet	1279
14.12.1.193. .spec.initContainers[].livenessProbe.httpGet.httpHeaders	1280
14.12.1.194. .spec.initContainers[].livenessProbe.httpGet.httpHeaders[]	1280
14.12.1.195. .spec.initContainers[].livenessProbe.tcpSocket	1281
14.12.1.196. .spec.initContainers[].ports	1281
14.12.1.197. .spec.initContainers[].ports[]	1281
14.12.1.198. .spec.initContainers[].readinessProbe	1282
14.12.1.199. .spec.initContainers[].readinessProbe.exec	1284
14.12.1.200. .spec.initContainers[].readinessProbe.grpc	1285
14.12.1.201. .spec.initContainers[].readinessProbe.httpGet	1285
14.12.1.202. .spec.initContainers[].readinessProbe.httpGet.httpHeaders	1286
14.12.1.203. .spec.initContainers[].readinessProbe.httpGet.httpHeaders[]	1286
14.12.1.204. .spec.initContainers[].readinessProbe.tcpSocket	1287
14.12.1.205. .spec.initContainers[].resizePolicy	1287
14.12.1.206. .spec.initContainers[].resizePolicy[]	1287
14.12.1.207. .spec.initContainers[].resources	1288
14.12.1.208. .spec.initContainers[].resources.claims	1289
14.12.1.209. .spec.initContainers[].resources.claims[]	1289
14.12.1.210. .spec.initContainers[].securityContext	1289
14.12.1.211. .spec.initContainers[].securityContext.capabilities	1292
14.12.1.212. .spec.initContainers[].securityContext.seLinuxOptions	1293
14.12.1.213. .spec.initContainers[].securityContext.seccompProfile	1293
14.12.1.214. .spec.initContainers[].securityContext.windowsOptions	1294
14.12.1.215. .spec.initContainers[].startupProbe	1295
14.12.1.216. .spec.initContainers[].startupProbe.exec	1297
14.12.1.217. .spec.initContainers[].startupProbe.grpc	1298
14.12.1.218. .spec.initContainers[].startupProbe.httpGet	1298
14.12.1.219. .spec.initContainers[].startupProbe.httpGet.httpHeaders	1299
14.12.1.220. .spec.initContainers[].startupProbe.httpGet.httpHeaders[]	1299
14.12.1.221. .spec.initContainers[].startupProbe.tcpSocket	1300
14.12.1.222. .spec.initContainers[].volumeDevices	1300
14.12.1.223. .spec.initContainers[].volumeDevices[]	1300
14.12.1.224. .spec.initContainers[].volumeMounts	1301
14.12.1.225. .spec.initContainers[].volumeMounts[]	1301
14.12.1.226. .spec.os	1303
14.12.1.227. .spec.readinessGates	1303
14.12.1.228. .spec.readinessGates[]	1304
14.12.1.229. .spec.resourceClaims	1304
14.12.1.230. .spec.resourceClaims[]	1304
14.12.1.231. .spec.resourceClaims[].source	1305
14.12.1.232. .spec.schedulingGates	1306
14.12.1.233. .spec.schedulingGates[]	1306
14.12.1.234. .spec.securityContext	1307
14.12.1.235. .spec.securityContext.seLinuxOptions	1310
14.12.1.236. .spec.securityContext.seccompProfile	1310
14.12.1.237. .spec.securityContext.sysctls	1311
14.12.1.238. .spec.securityContext.sysctls[]	1311
14.12.1.239. .spec.securityContext.windowsOptions	1312
14.12.1.240. .spec.tolerations	1313
14.12.1.241. .spec.tolerations[]	1313
14.12.1.242. .spec.topologySpreadConstraints	1315
14.12.1.243. .spec.topologySpreadConstraints[]	1315

14.12.1.244. .spec.volumes	1321
14.12.1.245. .spec.volumes[]	1321
14.12.1.246. .spec.volumes[].awsElasticBlockStore	1326
14.12.1.247. .spec.volumes[].azureDisk	1327
14.12.1.248. .spec.volumes[].azureFile	1328
14.12.1.249. .spec.volumes[].cephfs	1329
14.12.1.250. .spec.volumes[].cephfs.secretRef	1330
14.12.1.251. .spec.volumes[].cinder	1330
14.12.1.252. .spec.volumes[].cinder.secretRef	1331
14.12.1.253. .spec.volumes[].configMap	1331
14.12.1.254. .spec.volumes[].configMap.items	1332
14.12.1.255. .spec.volumes[].configMap.items[]	1333
14.12.1.256. .spec.volumes[].csi	1333
14.12.1.257. .spec.volumes[].csi.nodePublishSecretRef	1334
14.12.1.258. .spec.volumes[].downwardAPI	1335
14.12.1.259. .spec.volumes[].downwardAPI.items	1335
14.12.1.260. .spec.volumes[].downwardAPI.items[]	1335
14.12.1.261. .spec.volumes[].downwardAPI.items[].fieldRef	1336
14.12.1.262. .spec.volumes[].downwardAPI.items[].resourceFieldRef	1337
14.12.1.263. .spec.volumes[].emptyDir	1337
14.12.1.264. .spec.volumes[].ephemeral	1338
14.12.1.265. .spec.volumes[].ephemeral.volumeClaimTemplate	1338
14.12.1.266. .spec.volumes[].ephemeral.volumeClaimTemplate.spec	1339
14.12.1.267. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource	1341
14.12.1.268. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef	1342
14.12.1.269. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources	1343
14.12.1.270. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims	1344
14.12.1.271. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims[]	1344
14.12.1.272. .spec.volumes[].fc	1345
14.12.1.273. .spec.volumes[].flexVolume	1345
14.12.1.274. .spec.volumes[].flexVolume.secretRef	1346
14.12.1.275. .spec.volumes[].flocker	1346
14.12.1.276. .spec.volumes[].gcePersistentDisk	1347
14.12.1.277. .spec.volumes[].gitRepo	1348
14.12.1.278. .spec.volumes[].glusterfs	1349
14.12.1.279. .spec.volumes[].hostPath	1350
14.12.1.280. .spec.volumes[].iscsi	1351
14.12.1.281. .spec.volumes[].iscsi.secretRef	1353
14.12.1.282. .spec.volumes[].nfs	1353
14.12.1.283. .spec.volumes[].persistentVolumeClaim	1354
14.12.1.284. .spec.volumes[].photonPersistentDisk	1354
14.12.1.285. .spec.volumes[].portworxVolume	1355
14.12.1.286. .spec.volumes[].projected	1355
14.12.1.287. .spec.volumes[].projected.sources	1356
14.12.1.288. .spec.volumes[].projected.sources[]	1356
14.12.1.289. .spec.volumes[].projected.sources[].configMap	1357
14.12.1.290. .spec.volumes[].projected.sources[].configMap.items	1358
14.12.1.291. .spec.volumes[].projected.sources[].configMap.items[]	1358
14.12.1.292. .spec.volumes[].projected.sources[].downwardAPI	1359
14.12.1.293. .spec.volumes[].projected.sources[].downwardAPI.items	1359
14.12.1.294. .spec.volumes[].projected.sources[].downwardAPI.items[]	1360
14.12.1.295. .spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef	1360
14.12.1.296. .spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef	1361

14.12.1.297. .spec.volumes[].projected.sources[].secret	1361
14.12.1.298. .spec.volumes[].projected.sources[].secret.items	1362
14.12.1.299. .spec.volumes[].projected.sources[].secret.items[]	1362
14.12.1.300. .spec.volumes[].projected.sources[].serviceAccountToken	1363
14.12.1.301. .spec.volumes[].quobyte	1364
14.12.1.302. .spec.volumes[].rbd	1365
14.12.1.303. .spec.volumes[].rbd.secretRef	1366
14.12.1.304. .spec.volumes[].scaleIO	1367
14.12.1.305. .spec.volumes[].scaleIO.secretRef	1368
14.12.1.306. .spec.volumes[].secret	1368
14.12.1.307. .spec.volumes[].secret.items	1369
14.12.1.308. .spec.volumes[].secret.items[]	1370
14.12.1.309. .spec.volumes[].storageos	1370
14.12.1.310. .spec.volumes[].storageos.secretRef	1371
14.12.1.311. .spec.volumes[].vsphereVolume	1372
14.12.1.312. .status	1372
14.12.1.313. .status.conditions	1377
14.12.1.314. .status.conditions[]	1377
14.12.1.315. .status.containerStatuses	1378
14.12.1.316. .status.containerStatuses[]	1378
14.12.1.317. .status.containerStatuses[].lastState	1380
14.12.1.318. .status.containerStatuses[].lastState.running	1381
14.12.1.319. .status.containerStatuses[].lastState.terminated	1381
14.12.1.320. .status.containerStatuses[].lastState.waiting	1382
14.12.1.321. .status.containerStatuses[].resources	1382
14.12.1.322. .status.containerStatuses[].resources.claims	1383
14.12.1.323. .status.containerStatuses[].resources.claims[]	1383
14.12.1.324. .status.containerStatuses[].state	1384
14.12.1.325. .status.containerStatuses[].state.running	1384
14.12.1.326. .status.containerStatuses[].state.terminated	1384
14.12.1.327. .status.containerStatuses[].state.waiting	1385
14.12.1.328. .status.ephemeralContainerStatuses	1385
14.12.1.329. .status.ephemeralContainerStatuses[]	1386
14.12.1.330. .status.ephemeralContainerStatuses[].lastState	1388
14.12.1.331. .status.ephemeralContainerStatuses[].lastState.running	1388
14.12.1.332. .status.ephemeralContainerStatuses[].lastState.terminated	1389
14.12.1.333. .status.ephemeralContainerStatuses[].lastState.waiting	1389
14.12.1.334. .status.ephemeralContainerStatuses[].resources	1390
14.12.1.335. .status.ephemeralContainerStatuses[].resources.claims	1391
14.12.1.336. .status.ephemeralContainerStatuses[].resources.claims[]	1391
14.12.1.337. .status.ephemeralContainerStatuses[].state	1391
14.12.1.338. .status.ephemeralContainerStatuses[].state.running	1392
14.12.1.339. .status.ephemeralContainerStatuses[].state.terminated	1392
14.12.1.340. .status.ephemeralContainerStatuses[].state.waiting	1393
14.12.1.341. .status.initContainerStatuses	1393
14.12.1.342. .status.initContainerStatuses[]	1393
14.12.1.343. .status.initContainerStatuses[].lastState	1396
14.12.1.344. .status.initContainerStatuses[].lastState.running	1396
14.12.1.345. .status.initContainerStatuses[].lastState.terminated	1397
14.12.1.346. .status.initContainerStatuses[].lastState.waiting	1397
14.12.1.347. .status.initContainerStatuses[].resources	1398
14.12.1.348. .status.initContainerStatuses[].resources.claims	1399
14.12.1.349. .status.initContainerStatuses[].resources.claims[]	1399

14.12.1.350. .status.initContainerStatuses[].state	1399
14.12.1.351. .status.initContainerStatuses[].state.running	1400
14.12.1.352. .status.initContainerStatuses[].state.terminated	1400
14.12.1.353. .status.initContainerStatuses[].state.waiting	1401
14.12.1.354. .status.podIPs	1401
14.12.1.355. .status.podIPs[]	1401
14.12.2. API endpoints	1402
14.12.2.1. /api/v1/pods	1403
14.12.2.2. /api/v1/watch/pods	1406
14.12.2.3. /api/v1/namespaces/{namespace}/pods	1409
14.12.2.4. /api/v1/watch/namespaces/{namespace}/pods	1418
14.12.2.5. /api/v1/namespaces/{namespace}/pods/{name}	1422
14.12.2.6. /api/v1/namespaces/{namespace}/pods/{name}/log	1426
14.12.2.7. /api/v1/watch/namespaces/{namespace}/pods/{name}	1428
14.12.2.8. /api/v1/namespaces/{namespace}/pods/{name}/status	1432
14.12.2.9. /api/v1/namespaces/{namespace}/pods/{name}/ephemeralcontainers	1435
14.13. PODTEMPLATE [V1]	1438
14.13.1. Specification	1438
14.13.1.1. .template	1439
14.13.1.2. .template.spec	1440
14.13.1.3. .template.spec.affinity	1451
14.13.1.4. .template.spec.affinity.nodeAffinity	1451
14.13.1.5. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution	1452
14.13.1.6. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1452
14.13.1.7. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference	1453
14.13.1.8. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference	1453
14.13.1.9. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference	1454
14.13.1.10. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference	1454
14.13.1.11. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].preference	1454
14.13.1.12. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution	1455
14.13.1.13. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelector	1456
14.13.1.14. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelector	1456
14.13.1.15. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelector	1456
14.13.1.16. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelector	1457
14.13.1.17. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelector	1457
14.13.1.18. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelector	1458
14.13.1.19. .template.spec.affinity.podAffinity	1458
14.13.1.20. .template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution	1460
14.13.1.21. .template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1460
14.13.1.22. .template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAffinity	1461
14.13.1.23. .template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution	1462
14.13.1.24. .template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]	1462

14.13.1.25. .template.spec.affinity.podAntiAffinity	1463
14.13.1.26. .template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution	1465
14.13.1.27. .template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1465
14.13.1.28. .template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podAff	1466
14.13.1.29. .template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution	1467
14.13.1.30. .template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]	1467
14.13.1.31. .template.spec.containers	1468
14.13.1.32. .template.spec.containers[]	1468
14.13.1.33. .template.spec.containers[].env	1475
14.13.1.34. .template.spec.containers[].env[]	1475
14.13.1.35. .template.spec.containers[].env[].valueFrom	1476
14.13.1.36. .template.spec.containers[].env[].valueFrom.configMapKeyRef	1476
14.13.1.37. .template.spec.containers[].env[].valueFrom.fieldRef	1477
14.13.1.38. .template.spec.containers[].env[].valueFrom.resourceFieldRef	1477
14.13.1.39. .template.spec.containers[].env[].valueFrom.secretKeyRef	1478
14.13.1.40. .template.spec.containers[].envFrom	1478
14.13.1.41. .template.spec.containers[].envFrom[]	1478
14.13.1.42. .template.spec.containers[].envFrom[].configMapRef	1479
14.13.1.43. .template.spec.containers[].envFrom[].secretRef	1479
14.13.1.44. .template.spec.containers[].lifecycle	1480
14.13.1.45. .template.spec.containers[].lifecycle.postStart	1480
14.13.1.46. .template.spec.containers[].lifecycle.postStart.exec	1481
14.13.1.47. .template.spec.containers[].lifecycle.postStart.httpGet	1481
14.13.1.48. .template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders	1482
14.13.1.49. .template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]	1482
14.13.1.50. .template.spec.containers[].lifecycle.postStart.tcpSocket	1483
14.13.1.51. .template.spec.containers[].lifecycle.preStop	1483
14.13.1.52. .template.spec.containers[].lifecycle.preStop.exec	1484
14.13.1.53. .template.spec.containers[].lifecycle.preStop.httpGet	1484
14.13.1.54. .template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders	1485
14.13.1.55. .template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]	1485
14.13.1.56. .template.spec.containers[].lifecycle.preStop.tcpSocket	1486
14.13.1.57. .template.spec.containers[].livenessProbe	1486
14.13.1.58. .template.spec.containers[].livenessProbe.exec	1488
14.13.1.59. .template.spec.containers[].livenessProbe.grpc	1489
14.13.1.60. .template.spec.containers[].livenessProbe.httpGet	1489
14.13.1.61. .template.spec.containers[].livenessProbe.httpGet.httpHeaders	1490
14.13.1.62. .template.spec.containers[].livenessProbe.httpGet.httpHeaders[]	1490
14.13.1.63. .template.spec.containers[].livenessProbe.tcpSocket	1491
14.13.1.64. .template.spec.containers[].ports	1491
14.13.1.65. .template.spec.containers[].ports[]	1491
14.13.1.66. .template.spec.containers[].readinessProbe	1492
14.13.1.67. .template.spec.containers[].readinessProbe.exec	1494
14.13.1.68. .template.spec.containers[].readinessProbe.grpc	1495
14.13.1.69. .template.spec.containers[].readinessProbe.httpGet	1495
14.13.1.70. .template.spec.containers[].readinessProbe.httpGet.httpHeaders	1496
14.13.1.71. .template.spec.containers[].readinessProbe.httpGet.httpHeaders[]	1496
14.13.1.72. .template.spec.containers[].readinessProbe.tcpSocket	1497
14.13.1.73. .template.spec.containers[].resizePolicy	1497
14.13.1.74. .template.spec.containers[].resizePolicy[]	1497
14.13.1.75. .template.spec.containers[].resources	1498
14.13.1.76. .template.spec.containers[].resources.claims	1499

14.13.1.77. .template.spec.containers[].resources.claims[]	1499
14.13.1.78. .template.spec.containers[].securityContext	1499
14.13.1.79. .template.spec.containers[].securityContext.capabilities	1502
14.13.1.80. .template.spec.containers[].securityContext.seLinuxOptions	1503
14.13.1.81. .template.spec.containers[].securityContext.seccompProfile	1503
14.13.1.82. .template.spec.containers[].securityContext.windowsOptions	1504
14.13.1.83. .template.spec.containers[].startupProbe	1505
14.13.1.84. .template.spec.containers[].startupProbe.exec	1507
14.13.1.85. .template.spec.containers[].startupProbe.grpc	1508
14.13.1.86. .template.spec.containers[].startupProbe.httpGet	1508
14.13.1.87. .template.spec.containers[].startupProbe.httpGet.httpHeaders	1509
14.13.1.88. .template.spec.containers[].startupProbe.httpGet.httpHeaders[]	1509
14.13.1.89. .template.spec.containers[].startupProbe.tcpSocket	1510
14.13.1.90. .template.spec.containers[].volumeDevices	1510
14.13.1.91. .template.spec.containers[].volumeDevices[]	1510
14.13.1.92. .template.spec.containers[].volumeMounts	1511
14.13.1.93. .template.spec.containers[].volumeMounts[]	1511
14.13.1.94. .template.spec.dnsConfig	1513
14.13.1.95. .template.spec.dnsConfig.options	1513
14.13.1.96. .template.spec.dnsConfig.options[]	1514
14.13.1.97. .template.spec.ephemeralContainers	1514
14.13.1.98. .template.spec.ephemeralContainers[]	1514
14.13.1.99. .template.spec.ephemeralContainers[].env	1521
14.13.1.100. .template.spec.ephemeralContainers[].env[]	1521
14.13.1.101. .template.spec.ephemeralContainers[].env[].valueFrom	1522
14.13.1.102. .template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef	1522
14.13.1.103. .template.spec.ephemeralContainers[].env[].valueFrom.fieldRef	1523
14.13.1.104. .template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef	1523
14.13.1.105. .template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef	1524
14.13.1.106. .template.spec.ephemeralContainers[].envFrom	1524
14.13.1.107. .template.spec.ephemeralContainers[].envFrom[]	1524
14.13.1.108. .template.spec.ephemeralContainers[].envFrom[].configMapRef	1525
14.13.1.109. .template.spec.ephemeralContainers[].envFrom[].secretRef	1525
14.13.1.110. .template.spec.ephemeralContainers[].lifecycle	1526
14.13.1.111. .template.spec.ephemeralContainers[].lifecycle.postStart	1526
14.13.1.112. .template.spec.ephemeralContainers[].lifecycle.postStart.exec	1527
14.13.1.113. .template.spec.ephemeralContainers[].lifecycle.postStart.httpGet	1527
14.13.1.114. .template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders	1528
14.13.1.115. .template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]	1528
14.13.1.116. .template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket	1529
14.13.1.117. .template.spec.ephemeralContainers[].lifecycle.preStop	1529
14.13.1.118. .template.spec.ephemeralContainers[].lifecycle.preStop.exec	1530
14.13.1.119. .template.spec.ephemeralContainers[].lifecycle.preStop.httpGet	1530
14.13.1.120. .template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders	1531
14.13.1.121. .template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]	1531
14.13.1.122. .template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket	1532
14.13.1.123. .template.spec.ephemeralContainers[].livenessProbe	1532
14.13.1.124. .template.spec.ephemeralContainers[].livenessProbe.exec	1534
14.13.1.125. .template.spec.ephemeralContainers[].livenessProbe.grpc	1535
14.13.1.126. .template.spec.ephemeralContainers[].livenessProbe.httpGet	1535
14.13.1.127. .template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders	1536
14.13.1.128. .template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]	1536
14.13.1.129. .template.spec.ephemeralContainers[].livenessProbe.tcpSocket	1537



14.13.1.130. .template.spec.ephemeralContainers[].ports	1537
14.13.1.131. .template.spec.ephemeralContainers[].ports[]	1537
14.13.1.132. .template.spec.ephemeralContainers[].readinessProbe	1538
14.13.1.133. .template.spec.ephemeralContainers[].readinessProbe.exec	1540
14.13.1.134. .template.spec.ephemeralContainers[].readinessProbe.grpc	1541
14.13.1.135. .template.spec.ephemeralContainers[].readinessProbe.httpGet	1541
14.13.1.136. .template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders	1542
14.13.1.137. .template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]	1542
14.13.1.138. .template.spec.ephemeralContainers[].readinessProbe.tcpSocket	1543
14.13.1.139. .template.spec.ephemeralContainers[].resizePolicy	1543
14.13.1.140. .template.spec.ephemeralContainers[].resizePolicy[]	1543
14.13.1.141. .template.spec.ephemeralContainers[].resources	1544
14.13.1.142. .template.spec.ephemeralContainers[].resources.claims	1545
14.13.1.143. .template.spec.ephemeralContainers[].resources.claims[]	1545
14.13.1.144. .template.spec.ephemeralContainers[].securityContext	1545
14.13.1.145. .template.spec.ephemeralContainers[].securityContext.capabilities	1548
14.13.1.146. .template.spec.ephemeralContainers[].securityContext.seLinuxOptions	1549
14.13.1.147. .template.spec.ephemeralContainers[].securityContext.seccompProfile	1549
14.13.1.148. .template.spec.ephemeralContainers[].securityContext.windowsOptions	1550
14.13.1.149. .template.spec.ephemeralContainers[].startupProbe	1551
14.13.1.150. .template.spec.ephemeralContainers[].startupProbe.exec	1553
14.13.1.151. .template.spec.ephemeralContainers[].startupProbe.grpc	1554
14.13.1.152. .template.spec.ephemeralContainers[].startupProbe.httpGet	1554
14.13.1.153. .template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders	1555
14.13.1.154. .template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]	1555
14.13.1.155. .template.spec.ephemeralContainers[].startupProbe.tcpSocket	1556
14.13.1.156. .template.spec.ephemeralContainers[].volumeDevices	1556
14.13.1.157. .template.spec.ephemeralContainers[].volumeDevices[]	1556
14.13.1.158. .template.spec.ephemeralContainers[].volumeMounts	1557
14.13.1.159. .template.spec.ephemeralContainers[].volumeMounts[]	1557
14.13.1.160. .template.spec.hostAliases	1559
14.13.1.161. .template.spec.hostAliases[]	1559
14.13.1.162. .template.spec.imagePullSecrets	1559
14.13.1.163. .template.spec.imagePullSecrets[]	1559
14.13.1.164. .template.spec.initContainers	1560
14.13.1.165. .template.spec.initContainers[]	1560
14.13.1.166. .template.spec.initContainers[].env	1567
14.13.1.167. .template.spec.initContainers[].env[]	1567
14.13.1.168. .template.spec.initContainers[].env[].valueFrom	1568
14.13.1.169. .template.spec.initContainers[].env[].valueFrom.configMapKeyRef	1568
14.13.1.170. .template.spec.initContainers[].env[].valueFrom.fieldRef	1569
14.13.1.171. .template.spec.initContainers[].env[].valueFrom.resourceFieldRef	1569
14.13.1.172. .template.spec.initContainers[].env[].valueFrom.secretKeyRef	1570
14.13.1.173. .template.spec.initContainers[].envFrom	1570
14.13.1.174. .template.spec.initContainers[].envFrom[]	1570
14.13.1.175. .template.spec.initContainers[].envFrom[].configMapRef	1571
14.13.1.176. .template.spec.initContainers[].envFrom[].secretRef	1572
14.13.1.177. .template.spec.initContainers[].lifecycle	1572
14.13.1.178. .template.spec.initContainers[].lifecycle.postStart	1572
14.13.1.179. .template.spec.initContainers[].lifecycle.postStart.exec	1573
14.13.1.180. .template.spec.initContainers[].lifecycle.postStart.httpGet	1573
14.13.1.181. .template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders	1574
14.13.1.182. .template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]	1574

14.13.1.183. .template.spec.initContainers[].lifecycle.postStart.tcpSocket	1575
14.13.1.184. .template.spec.initContainers[].lifecycle.preStop	1575
14.13.1.185. .template.spec.initContainers[].lifecycle.preStop.exec	1576
14.13.1.186. .template.spec.initContainers[].lifecycle.preStop.httpGet	1576
14.13.1.187. .template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders	1577
14.13.1.188. .template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]	1577
14.13.1.189. .template.spec.initContainers[].lifecycle.preStop.tcpSocket	1578
14.13.1.190. .template.spec.initContainers[].livenessProbe	1578
14.13.1.191. .template.spec.initContainers[].livenessProbe.exec	1580
14.13.1.192. .template.spec.initContainers[].livenessProbe.grpc	1581
14.13.1.193. .template.spec.initContainers[].livenessProbe.httpGet	1581
14.13.1.194. .template.spec.initContainers[].livenessProbe.httpGet.httpHeaders	1582
14.13.1.195. .template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]	1582
14.13.1.196. .template.spec.initContainers[].livenessProbe.tcpSocket	1583
14.13.1.197. .template.spec.initContainers[].ports	1583
14.13.1.198. .template.spec.initContainers[].ports[]	1583
14.13.1.199. .template.spec.initContainers[].readinessProbe	1584
14.13.1.200. .template.spec.initContainers[].readinessProbe.exec	1586
14.13.1.201. .template.spec.initContainers[].readinessProbe.grpc	1587
14.13.1.202. .template.spec.initContainers[].readinessProbe.httpGet	1587
14.13.1.203. .template.spec.initContainers[].readinessProbe.httpGet.httpHeaders	1588
14.13.1.204. .template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]	1588
14.13.1.205. .template.spec.initContainers[].readinessProbe.tcpSocket	1589
14.13.1.206. .template.spec.initContainers[].resizePolicy	1589
14.13.1.207. .template.spec.initContainers[].resizePolicy[]	1589
14.13.1.208. .template.spec.initContainers[].resources	1590
14.13.1.209. .template.spec.initContainers[].resources.claims	1591
14.13.1.210. .template.spec.initContainers[].resources.claims[]	1591
14.13.1.211. .template.spec.initContainers[].securityContext	1591
14.13.1.212. .template.spec.initContainers[].securityContext.capabilities	1594
14.13.1.213. .template.spec.initContainers[].securityContext.seLinuxOptions	1595
14.13.1.214. .template.spec.initContainers[].securityContext.seccompProfile	1595
14.13.1.215. .template.spec.initContainers[].securityContext.windowsOptions	1596
14.13.1.216. .template.spec.initContainers[].startupProbe	1597
14.13.1.217. .template.spec.initContainers[].startupProbe.exec	1599
14.13.1.218. .template.spec.initContainers[].startupProbe.grpc	1600
14.13.1.219. .template.spec.initContainers[].startupProbe.httpGet	1600
14.13.1.220. .template.spec.initContainers[].startupProbe.httpGet.httpHeaders	1601
14.13.1.221. .template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]	1601
14.13.1.222. .template.spec.initContainers[].startupProbe.tcpSocket	1602
14.13.1.223. .template.spec.initContainers[].volumeDevices	1602
14.13.1.224. .template.spec.initContainers[].volumeDevices[]	1602
14.13.1.225. .template.spec.initContainers[].volumeMounts	1603
14.13.1.226. .template.spec.initContainers[].volumeMounts[]	1603
14.13.1.227. .template.spec.os	1605
14.13.1.228. .template.spec.readinessGates	1605
14.13.1.229. .template.spec.readinessGates[]	1606
14.13.1.230. .template.spec.resourceClaims	1606
14.13.1.231. .template.spec.resourceClaims[]	1606
14.13.1.232. .template.spec.resourceClaims[].source	1607
14.13.1.233. .template.spec.schedulingGates	1608
14.13.1.234. .template.spec.schedulingGates[]	1608
14.13.1.235. .template.spec.securityContext	1609

14.13.1.236. .template.spec.securityContext.seLinuxOptions	1612
14.13.1.237. .template.spec.securityContext.seccompProfile	1612
14.13.1.238. .template.spec.securityContext.sysctls	1613
14.13.1.239. .template.spec.securityContext.sysctls[]	1613
14.13.1.240. .template.spec.securityContext.windowsOptions	1614
14.13.1.241. .template.spec.tolerations	1615
14.13.1.242. .template.spec.tolerations[]	1615
14.13.1.243. .template.spec.topologySpreadConstraints	1617
14.13.1.244. .template.spec.topologySpreadConstraints[]	1617
14.13.1.245. .template.spec.volumes	1623
14.13.1.246. .template.spec.volumes[]	1623
14.13.1.247. .template.spec.volumes[].awsElasticBlockStore	1628
14.13.1.248. .template.spec.volumes[].azureDisk	1629
14.13.1.249. .template.spec.volumes[].azureFile	1630
14.13.1.250. .template.spec.volumes[].cephfs	1631
14.13.1.251. .template.spec.volumes[].cephfs.secretRef	1632
14.13.1.252. .template.spec.volumes[].cinder	1632
14.13.1.253. .template.spec.volumes[].cinder.secretRef	1633
14.13.1.254. .template.spec.volumes[].configMap	1633
14.13.1.255. .template.spec.volumes[].configMap.items	1634
14.13.1.256. .template.spec.volumes[].configMap.items[]	1635
14.13.1.257. .template.spec.volumes[].csi	1635
14.13.1.258. .template.spec.volumes[].csi.nodePublishSecretRef	1636
14.13.1.259. .template.spec.volumes[].downwardAPI	1637
14.13.1.260. .template.spec.volumes[].downwardAPI.items	1637
14.13.1.261. .template.spec.volumes[].downwardAPI.items[]	1637
14.13.1.262. .template.spec.volumes[].downwardAPI.items[].fieldRef	1638
14.13.1.263. .template.spec.volumes[].downwardAPI.items[].resourceFieldRef	1639
14.13.1.264. .template.spec.volumes[].emptyDir	1639
14.13.1.265. .template.spec.volumes[].ephemeral	1640
14.13.1.266. .template.spec.volumes[].ephemeral.volumeClaimTemplate	1640
14.13.1.267. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec	1641
14.13.1.268. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource	1643
14.13.1.269. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef	1644
14.13.1.270. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources	1645
14.13.1.271. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims	1646
14.13.1.272. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims[]	1646
14.13.1.273. .template.spec.volumes[].fc	1647
14.13.1.274. .template.spec.volumes[].flexVolume	1647
14.13.1.275. .template.spec.volumes[].flexVolume.secretRef	1648
14.13.1.276. .template.spec.volumes[].flocker	1648
14.13.1.277. .template.spec.volumes[].gcePersistentDisk	1649
14.13.1.278. .template.spec.volumes[].gitRepo	1650
14.13.1.279. .template.spec.volumes[].glusterfs	1651
14.13.1.280. .template.spec.volumes[].hostPath	1652
14.13.1.281. .template.spec.volumes[].iscsi	1653
14.13.1.282. .template.spec.volumes[].iscsi.secretRef	1655
14.13.1.283. .template.spec.volumes[].nfs	1655
14.13.1.284. .template.spec.volumes[].persistentVolumeClaim	1656
14.13.1.285. .template.spec.volumes[].photonPersistentDisk	1656
14.13.1.286. .template.spec.volumes[].portworxVolume	1657
14.13.1.287. .template.spec.volumes[].projected	1657
14.13.1.288. .template.spec.volumes[].projected.sources	1658

14.13.1.289. .template.spec.volumes[].projected.sources[]	1658
14.13.1.290. .template.spec.volumes[].projected.sources[].configMap	1659
14.13.1.291. .template.spec.volumes[].projected.sources[].configMap.items	1660
14.13.1.292. .template.spec.volumes[].projected.sources[].configMap.items[]	1660
14.13.1.293. .template.spec.volumes[].projected.sources[].downwardAPI	1661
14.13.1.294. .template.spec.volumes[].projected.sources[].downwardAPI.items	1661
14.13.1.295. .template.spec.volumes[].projected.sources[].downwardAPI.items[]	1662
14.13.1.296. .template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef	1662
14.13.1.297. .template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef	1663
14.13.1.298. .template.spec.volumes[].projected.sources[].secret	1663
14.13.1.299. .template.spec.volumes[].projected.sources[].secret.items	1664
14.13.1.300. .template.spec.volumes[].projected.sources[].secret.items[]	1664
14.13.1.301. .template.spec.volumes[].projected.sources[].serviceAccountToken	1665
14.13.1.302. .template.spec.volumes[].quobyte	1666
14.13.1.303. .template.spec.volumes[].rbd	1667
14.13.1.304. .template.spec.volumes[].rbd.secretRef	1668
14.13.1.305. .template.spec.volumes[].scaleIO	1669
14.13.1.306. .template.spec.volumes[].scaleIO.secretRef	1670
14.13.1.307. .template.spec.volumes[].secret	1670
14.13.1.308. .template.spec.volumes[].secret.items	1671
14.13.1.309. .template.spec.volumes[].secret.items[]	1672
14.13.1.310. .template.spec.volumes[].storageos	1672
14.13.1.311. .template.spec.volumes[].storageos.secretRef	1673
14.13.1.312. .template.spec.volumes[].vsphereVolume	1674
14.13.2. API endpoints	1674
14.13.2.1. /api/v1/podtemplates	1675
14.13.2.2. /api/v1/watch/podtemplates	1679
14.13.2.3. /api/v1/namespaces/{namespace}/podtemplates	1682
14.13.2.4. /api/v1/watch/namespaces/{namespace}/podtemplates	1691
14.13.2.5. /api/v1/namespaces/{namespace}/podtemplates/{name}	1695
14.13.2.6. /api/v1/watch/namespaces/{namespace}/podtemplates/{name}	1699
14.14. REPLICATIONCONTROLLER [V1]	1703
14.14.1. Specification	1703
14.14.1.1. .spec	1704
14.14.1.2. .spec.template	1705
14.14.1.3. .spec.template.spec	1705
14.14.1.4. .spec.template.spec.affinity	1717
14.14.1.5. .spec.template.spec.affinity.nodeAffinity	1717
14.14.1.6. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution	1718
14.14.1.7. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1718
14.14.1.8. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].prefe	1719
14.14.1.9. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].prefe	1719
14.14.1.10. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].prefe	1720
14.14.1.11. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].prefe	1720
14.14.1.12. .spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].prefe	1720
14.14.1.13. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution	1721
14.14.1.14. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSe	

	1722
14.14.1.15. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSe	1722
14.14.1.16. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSe	1722
14.14.1.17. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSe	1723
14.14.1.18. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSe	1723
14.14.1.19. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSe	1724
14.14.1.20. .spec.template.spec.affinity.podAffinity	1724
14.14.1.21. .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution	1726
14.14.1.22. .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1726
14.14.1.23. .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].podA	1727
14.14.1.24. .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution	1728
14.14.1.25. .spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]	1728
14.14.1.26. .spec.template.spec.affinity.podAntiAffinity	1729
14.14.1.27. .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution	1731
14.14.1.28. .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]	1731
14.14.1.29. .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[].p	1732
14.14.1.30. .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution	1733
14.14.1.31. .spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]	1733
14.14.1.32. .spec.template.spec.containers	1734
14.14.1.33. .spec.template.spec.containers[]	1734
14.14.1.34. .spec.template.spec.containers[].env	1741
14.14.1.35. .spec.template.spec.containers[].env[]	1741
14.14.1.36. .spec.template.spec.containers[].env[].valueFrom	1742
14.14.1.37. .spec.template.spec.containers[].env[].valueFrom.configMapKeyRef	1742
14.14.1.38. .spec.template.spec.containers[].env[].valueFrom.fieldRef	1743
14.14.1.39. .spec.template.spec.containers[].env[].valueFrom.resourceFieldRef	1743
14.14.1.40. .spec.template.spec.containers[].env[].valueFrom.secretKeyRef	1744
14.14.1.41. .spec.template.spec.containers[].envFrom	1744
14.14.1.42. .spec.template.spec.containers[].envFrom[]	1744
14.14.1.43. .spec.template.spec.containers[].envFrom[].configMapRef	1745
14.14.1.44. .spec.template.spec.containers[].envFrom[].secretRef	1745
14.14.1.45. .spec.template.spec.containers[].lifecycle	1746
14.14.1.46. .spec.template.spec.containers[].lifecycle.postStart	1746
14.14.1.47. .spec.template.spec.containers[].lifecycle.postStart.exec	1747
14.14.1.48. .spec.template.spec.containers[].lifecycle.postStart.httpGet	1747
14.14.1.49. .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders	1748
14.14.1.50. .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]	1748
14.14.1.51. .spec.template.spec.containers[].lifecycle.postStart.tcpSocket	1749
14.14.1.52. .spec.template.spec.containers[].lifecycle.preStop	1749
14.14.1.53. .spec.template.spec.containers[].lifecycle.preStop.exec	1750
14.14.1.54. .spec.template.spec.containers[].lifecycle.preStop.httpGet	1750
14.14.1.55. .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders	1751

14.14.1.56. .spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]	1751
14.14.1.57. .spec.template.spec.containers[].lifecycle.preStop.tcpSocket	1752
14.14.1.58. .spec.template.spec.containers[].livenessProbe	1752
14.14.1.59. .spec.template.spec.containers[].livenessProbe.exec	1754
14.14.1.60. .spec.template.spec.containers[].livenessProbe.grpc	1755
14.14.1.61. .spec.template.spec.containers[].livenessProbe.httpGet	1755
14.14.1.62. .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders	1756
14.14.1.63. .spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]	1756
14.14.1.64. .spec.template.spec.containers[].livenessProbe.tcpSocket	1757
14.14.1.65. .spec.template.spec.containers[].ports	1757
14.14.1.66. .spec.template.spec.containers[].ports[]	1757
14.14.1.67. .spec.template.spec.containers[].readinessProbe	1758
14.14.1.68. .spec.template.spec.containers[].readinessProbe.exec	1760
14.14.1.69. .spec.template.spec.containers[].readinessProbe.grpc	1761
14.14.1.70. .spec.template.spec.containers[].readinessProbe.httpGet	1761
14.14.1.71. .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders	1762
14.14.1.72. .spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]	1762
14.14.1.73. .spec.template.spec.containers[].readinessProbe.tcpSocket	1763
14.14.1.74. .spec.template.spec.containers[].resizePolicy	1763
14.14.1.75. .spec.template.spec.containers[].resizePolicy[]	1763
14.14.1.76. .spec.template.spec.containers[].resources	1764
14.14.1.77. .spec.template.spec.containers[].resources.claims	1765
14.14.1.78. .spec.template.spec.containers[].resources.claims[]	1765
14.14.1.79. .spec.template.spec.containers[].securityContext	1765
14.14.1.80. .spec.template.spec.containers[].securityContext.capabilities	1768
14.14.1.81. .spec.template.spec.containers[].securityContext.seLinuxOptions	1769
14.14.1.82. .spec.template.spec.containers[].securityContext.seccompProfile	1769
14.14.1.83. .spec.template.spec.containers[].securityContext.windowsOptions	1770
14.14.1.84. .spec.template.spec.containers[].startupProbe	1771
14.14.1.85. .spec.template.spec.containers[].startupProbe.exec	1773
14.14.1.86. .spec.template.spec.containers[].startupProbe.grpc	1774
14.14.1.87. .spec.template.spec.containers[].startupProbe.httpGet	1774
14.14.1.88. .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders	1775
14.14.1.89. .spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]	1775
14.14.1.90. .spec.template.spec.containers[].startupProbe.tcpSocket	1776
14.14.1.91. .spec.template.spec.containers[].volumeDevices	1776
14.14.1.92. .spec.template.spec.containers[].volumeDevices[]	1776
14.14.1.93. .spec.template.spec.containers[].volumeMounts	1777
14.14.1.94. .spec.template.spec.containers[].volumeMounts[]	1777
14.14.1.95. .spec.template.spec.dnsConfig	1779
14.14.1.96. .spec.template.spec.dnsConfig.options	1779
14.14.1.97. .spec.template.spec.dnsConfig.options[]	1780
14.14.1.98. .spec.template.spec.ephemeralContainers	1780
14.14.1.99. .spec.template.spec.ephemeralContainers[]	1780
14.14.1.100. .spec.template.spec.ephemeralContainers[].env	1787
14.14.1.101. .spec.template.spec.ephemeralContainers[].env[]	1787
14.14.1.102. .spec.template.spec.ephemeralContainers[].env[].valueFrom	1788
14.14.1.103. .spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef	1788
14.14.1.104. .spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef	1789
14.14.1.105. .spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef	1789
14.14.1.106. .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef	1790
14.14.1.107. .spec.template.spec.ephemeralContainers[].envFrom	1790
14.14.1.108. .spec.template.spec.ephemeralContainers[].envFrom[]	1790

14.14.1.109. .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef	1791
14.14.1.110. .spec.template.spec.ephemeralContainers[].envFrom[].secretRef	1791
14.14.1.111. .spec.template.spec.ephemeralContainers[].lifecycle	1792
14.14.1.112. .spec.template.spec.ephemeralContainers[].lifecycle.postStart	1792
14.14.1.113. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec	1793
14.14.1.114. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet	1793
14.14.1.115. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders	1794
14.14.1.116. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]	1794
14.14.1.117. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket	1795
14.14.1.118. .spec.template.spec.ephemeralContainers[].lifecycle.preStop	1795
14.14.1.119. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec	1796
14.14.1.120. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet	1796
14.14.1.121. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders	1797
14.14.1.122. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]	1797
14.14.1.123. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket	1798
14.14.1.124. .spec.template.spec.ephemeralContainers[].livenessProbe	1798
14.14.1.125. .spec.template.spec.ephemeralContainers[].livenessProbe.exec	1800
14.14.1.126. .spec.template.spec.ephemeralContainers[].livenessProbe.grpc	1801
14.14.1.127. .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet	1801
14.14.1.128. .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders	1802
14.14.1.129. .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]	1802
14.14.1.130. .spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket	1803
14.14.1.131. .spec.template.spec.ephemeralContainers[].ports	1803
14.14.1.132. .spec.template.spec.ephemeralContainers[].ports[]	1803
14.14.1.133. .spec.template.spec.ephemeralContainers[].readinessProbe	1804
14.14.1.134. .spec.template.spec.ephemeralContainers[].readinessProbe.exec	1806
14.14.1.135. .spec.template.spec.ephemeralContainers[].readinessProbe.grpc	1807
14.14.1.136. .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet	1807
14.14.1.137. .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders	1808
14.14.1.138. .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]	1808
14.14.1.139. .spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket	1809
14.14.1.140. .spec.template.spec.ephemeralContainers[].resizePolicy	1809
14.14.1.141. .spec.template.spec.ephemeralContainers[].resizePolicy[]	1809
14.14.1.142. .spec.template.spec.ephemeralContainers[].resources	1810
14.14.1.143. .spec.template.spec.ephemeralContainers[].resources.claims	1811
14.14.1.144. .spec.template.spec.ephemeralContainers[].resources.claims[]	1811
14.14.1.145. .spec.template.spec.ephemeralContainers[].securityContext	1811
14.14.1.146. .spec.template.spec.ephemeralContainers[].securityContext.capabilities	1814
14.14.1.147. .spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions	1815
14.14.1.148. .spec.template.spec.ephemeralContainers[].securityContext.seccompProfile	1815
14.14.1.149. .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions	1816
14.14.1.150. .spec.template.spec.ephemeralContainers[].startupProbe	1817
14.14.1.151. .spec.template.spec.ephemeralContainers[].startupProbe.exec	1819
14.14.1.152. .spec.template.spec.ephemeralContainers[].startupProbe.grpc	1820
14.14.1.153. .spec.template.spec.ephemeralContainers[].startupProbe.httpGet	1820
14.14.1.154. .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders	1821
14.14.1.155. .spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]	1821
14.14.1.156. .spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket	1822
14.14.1.157. .spec.template.spec.ephemeralContainers[].volumeDevices	1822
14.14.1.158. .spec.template.spec.ephemeralContainers[].volumeDevices[]	1822
14.14.1.159. .spec.template.spec.ephemeralContainers[].volumeMounts	1823
14.14.1.160. .spec.template.spec.ephemeralContainers[].volumeMounts[]	1823
14.14.1.161. .spec.template.spec.hostAliases	1825

14.14.1.162. .spec.template.spec.hostAliases[]	1825
14.14.1.163. .spec.template.spec.imagePullSecrets	1825
14.14.1.164. .spec.template.spec.imagePullSecrets[]	1825
14.14.1.165. .spec.template.spec.initContainers	1826
14.14.1.166. .spec.template.spec.initContainers[]	1826
14.14.1.167. .spec.template.spec.initContainers[].env	1833
14.14.1.168. .spec.template.spec.initContainers[].env[]	1833
14.14.1.169. .spec.template.spec.initContainers[].env[].valueFrom	1834
14.14.1.170. .spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef	1834
14.14.1.171. .spec.template.spec.initContainers[].env[].valueFrom.fieldRef	1835
14.14.1.172. .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef	1835
14.14.1.173. .spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef	1836
14.14.1.174. .spec.template.spec.initContainers[].envFrom	1836
14.14.1.175. .spec.template.spec.initContainers[].envFrom[]	1836
14.14.1.176. .spec.template.spec.initContainers[].envFrom[].configMapRef	1837
14.14.1.177. .spec.template.spec.initContainers[].envFrom[].secretRef	1838
14.14.1.178. .spec.template.spec.initContainers[].lifecycle	1838
14.14.1.179. .spec.template.spec.initContainers[].lifecycle.postStart	1838
14.14.1.180. .spec.template.spec.initContainers[].lifecycle.postStart.exec	1839
14.14.1.181. .spec.template.spec.initContainers[].lifecycle.postStart.httpGet	1839
14.14.1.182. .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders	1840
14.14.1.183. .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]	1840
14.14.1.184. .spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket	1841
14.14.1.185. .spec.template.spec.initContainers[].lifecycle.preStop	1841
14.14.1.186. .spec.template.spec.initContainers[].lifecycle.preStop.exec	1842
14.14.1.187. .spec.template.spec.initContainers[].lifecycle.preStop.httpGet	1842
14.14.1.188. .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders	1843
14.14.1.189. .spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]	1843
14.14.1.190. .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket	1844
14.14.1.191. .spec.template.spec.initContainers[].livenessProbe	1844
14.14.1.192. .spec.template.spec.initContainers[].livenessProbe.exec	1846
14.14.1.193. .spec.template.spec.initContainers[].livenessProbe.grpc	1847
14.14.1.194. .spec.template.spec.initContainers[].livenessProbe.httpGet	1847
14.14.1.195. .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders	1848
14.14.1.196. .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]	1848
14.14.1.197. .spec.template.spec.initContainers[].livenessProbe.tcpSocket	1849
14.14.1.198. .spec.template.spec.initContainers[].ports	1849
14.14.1.199. .spec.template.spec.initContainers[].ports[]	1849
14.14.1.200. .spec.template.spec.initContainers[].readinessProbe	1850
14.14.1.201. .spec.template.spec.initContainers[].readinessProbe.exec	1852
14.14.1.202. .spec.template.spec.initContainers[].readinessProbe.grpc	1853
14.14.1.203. .spec.template.spec.initContainers[].readinessProbe.httpGet	1853
14.14.1.204. .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders	1854
14.14.1.205. .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]	1854
14.14.1.206. .spec.template.spec.initContainers[].readinessProbe.tcpSocket	1855
14.14.1.207. .spec.template.spec.initContainers[].resizePolicy	1855
14.14.1.208. .spec.template.spec.initContainers[].resizePolicy[]	1855
14.14.1.209. .spec.template.spec.initContainers[].resources	1856
14.14.1.210. .spec.template.spec.initContainers[].resources.claims	1857
14.14.1.211. .spec.template.spec.initContainers[].resources.claims[]	1857
14.14.1.212. .spec.template.spec.initContainers[].securityContext	1857
14.14.1.213. .spec.template.spec.initContainers[].securityContext.capabilities	1860
14.14.1.214. .spec.template.spec.initContainers[].securityContext.seLinuxOptions	1861



14.14.1.215. .spec.template.spec.initContainers[].securityContext.seccompProfile	1861
14.14.1.216. .spec.template.spec.initContainers[].securityContext.windowsOptions	1862
14.14.1.217. .spec.template.spec.initContainers[].startupProbe	1863
14.14.1.218. .spec.template.spec.initContainers[].startupProbe.exec	1865
14.14.1.219. .spec.template.spec.initContainers[].startupProbe.grpc	1866
14.14.1.220. .spec.template.spec.initContainers[].startupProbe.httpGet	1866
14.14.1.221. .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders	1867
14.14.1.222. .spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]	1867
14.14.1.223. .spec.template.spec.initContainers[].startupProbe.tcpSocket	1868
14.14.1.224. .spec.template.spec.initContainers[].volumeDevices	1868
14.14.1.225. .spec.template.spec.initContainers[].volumeDevices[]	1868
14.14.1.226. .spec.template.spec.initContainers[].volumeMounts	1869
14.14.1.227. .spec.template.spec.initContainers[].volumeMounts[]	1869
14.14.1.228. .spec.template.spec.os	1871
14.14.1.229. .spec.template.spec.readinessGates	1871
14.14.1.230. .spec.template.spec.readinessGates[]	1872
14.14.1.231. .spec.template.spec.resourceClaims	1872
14.14.1.232. .spec.template.spec.resourceClaims[]	1872
14.14.1.233. .spec.template.spec.resourceClaims[].source	1873
14.14.1.234. .spec.template.spec.schedulingGates	1874
14.14.1.235. .spec.template.spec.schedulingGates[]	1874
14.14.1.236. .spec.template.spec.securityContext	1875
14.14.1.237. .spec.template.spec.securityContext.seLinuxOptions	1878
14.14.1.238. .spec.template.spec.securityContext.seccompProfile	1878
14.14.1.239. .spec.template.spec.securityContext.sysctls	1879
14.14.1.240. .spec.template.spec.securityContext.sysctls[]	1879
14.14.1.241. .spec.template.spec.securityContext.windowsOptions	1880
14.14.1.242. .spec.template.spec.tolerations	1881
14.14.1.243. .spec.template.spec.tolerations[]	1881
14.14.1.244. .spec.template.spec.topologySpreadConstraints	1883
14.14.1.245. .spec.template.spec.topologySpreadConstraints[]	1883
14.14.1.246. .spec.template.spec.volumes	1889
14.14.1.247. .spec.template.spec.volumes[]	1889
14.14.1.248. .spec.template.spec.volumes[].awsElasticBlockStore	1894
14.14.1.249. .spec.template.spec.volumes[].azureDisk	1895
14.14.1.250. .spec.template.spec.volumes[].azureFile	1896
14.14.1.251. .spec.template.spec.volumes[].cephfs	1897
14.14.1.252. .spec.template.spec.volumes[].cephfs.secretRef	1898
14.14.1.253. .spec.template.spec.volumes[].cinder	1898
14.14.1.254. .spec.template.spec.volumes[].cinder.secretRef	1899
14.14.1.255. .spec.template.spec.volumes[].configMap	1899
14.14.1.256. .spec.template.spec.volumes[].configMap.items	1900
14.14.1.257. .spec.template.spec.volumes[].configMap.items[]	1901
14.14.1.258. .spec.template.spec.volumes[].csi	1901
14.14.1.259. .spec.template.spec.volumes[].csi.nodePublishSecretRef	1902
14.14.1.260. .spec.template.spec.volumes[].downwardAPI	1903
14.14.1.261. .spec.template.spec.volumes[].downwardAPI.items	1903
14.14.1.262. .spec.template.spec.volumes[].downwardAPI.items[]	1903
14.14.1.263. .spec.template.spec.volumes[].downwardAPI.items[].fieldRef	1904
14.14.1.264. .spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef	1905
14.14.1.265. .spec.template.spec.volumes[].emptyDir	1905
14.14.1.266. .spec.template.spec.volumes[].ephemeral	1906
14.14.1.267. .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate	1906

14.14.1.268.	.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec	1907
14.14.1.269.	.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource	1909
14.14.1.270.	.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef	1910
14.14.1.271.	.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources	1911
14.14.1.272.	.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims	1912
14.14.1.273.	.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims[]	1912
14.14.1.274.	.spec.template.spec.volumes[].fc	1913
14.14.1.275.	.spec.template.spec.volumes[].flexVolume	1913
14.14.1.276.	.spec.template.spec.volumes[].flexVolume.secretRef	1914
14.14.1.277.	.spec.template.spec.volumes[].flocker	1914
14.14.1.278.	.spec.template.spec.volumes[].gcePersistentDisk	1915
14.14.1.279.	.spec.template.spec.volumes[].gitRepo	1916
14.14.1.280.	.spec.template.spec.volumes[].glusterfs	1917
14.14.1.281.	.spec.template.spec.volumes[].hostPath	1918
14.14.1.282.	.spec.template.spec.volumes[].iscsi	1919
14.14.1.283.	.spec.template.spec.volumes[].iscsi.secretRef	1921
14.14.1.284.	.spec.template.spec.volumes[].nfs	1921
14.14.1.285.	.spec.template.spec.volumes[].persistentVolumeClaim	1922
14.14.1.286.	.spec.template.spec.volumes[].photonPersistentDisk	1922
14.14.1.287.	.spec.template.spec.volumes[].portworxVolume	1923
14.14.1.288.	.spec.template.spec.volumes[].projected	1923
14.14.1.289.	.spec.template.spec.volumes[].projected.sources	1924
14.14.1.290.	.spec.template.spec.volumes[].projected.sources[]	1924
14.14.1.291.	.spec.template.spec.volumes[].projected.sources[].configMap	1925
14.14.1.292.	.spec.template.spec.volumes[].projected.sources[].configMap.items	1926
14.14.1.293.	.spec.template.spec.volumes[].projected.sources[].configMap.items[]	1926
14.14.1.294.	.spec.template.spec.volumes[].projected.sources[].downwardAPI	1927
14.14.1.295.	.spec.template.spec.volumes[].projected.sources[].downwardAPI.items	1927
14.14.1.296.	.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]	1928
14.14.1.297.	.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef	1928
14.14.1.298.	.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef	1929
14.14.1.299.	.spec.template.spec.volumes[].projected.sources[].secret	1929
14.14.1.300.	.spec.template.spec.volumes[].projected.sources[].secret.items	1930
14.14.1.301.	.spec.template.spec.volumes[].projected.sources[].secret.items[]	1930
14.14.1.302.	.spec.template.spec.volumes[].projected.sources[].serviceAccountToken	1931
14.14.1.303.	.spec.template.spec.volumes[].quobyte	1932
14.14.1.304.	.spec.template.spec.volumes[].rbd	1933
14.14.1.305.	.spec.template.spec.volumes[].rbd.secretRef	1934
14.14.1.306.	.spec.template.spec.volumes[].scaleIO	1935
14.14.1.307.	.spec.template.spec.volumes[].scaleIO.secretRef	1936
14.14.1.308.	.spec.template.spec.volumes[].secret	1936
14.14.1.309.	.spec.template.spec.volumes[].secret.items	1937
14.14.1.310.	.spec.template.spec.volumes[].secret.items[]	1938
14.14.1.311.	.spec.template.spec.volumes[].storageos	1938
14.14.1.312.	.spec.template.spec.volumes[].storageos.secretRef	1939
14.14.1.313.	.spec.template.spec.volumes[].vsphereVolume	1940
14.14.1.314.	.status	1940
14.14.1.315.	.status.conditions	1941
14.14.1.316.	.status.conditions[]	1941
14.14.2.	API endpoints	1942
14.14.2.1.	/api/v1/replicationcontrollers	1943
14.14.2.2.	/api/v1/watch/replicationcontrollers	1947
14.14.2.3.	/api/v1/namespaces/{namespace}/replicationcontrollers	1950

14.14.2.4. /api/v1/watch/namespaces/{namespace}/replicationcontrollers	1959
14.14.2.5. /api/v1/namespaces/{namespace}/replicationcontrollers/{name}	1963
14.14.2.6. /api/v1/watch/namespaces/{namespace}/replicationcontrollers/{name}	1967
14.14.2.7. /api/v1/namespaces/{namespace}/replicationcontrollers/{name}/status	1971
14.15. RESOURCEQUOTA [V1]	1974
14.15.1. Specification	1974
14.15.1.1. .spec	1975
14.15.1.2. .spec.scopeSelector	1975
14.15.1.3. .spec.scopeSelector.matchExpressions	1976
14.15.1.4. .spec.scopeSelector.matchExpressions[]	1976
14.15.1.5. .status	1977
14.15.2. API endpoints	1978
14.15.2.1. /api/v1/resourcequotas	1979
14.15.2.2. /api/v1/watch/resourcequotas	1982
14.15.2.3. /api/v1/namespaces/{namespace}/resourcequotas	1985
14.15.2.4. /api/v1/watch/namespaces/{namespace}/resourcequotas	1994
14.15.2.5. /api/v1/namespaces/{namespace}/resourcequotas/{name}	1998
14.15.2.6. /api/v1/watch/namespaces/{namespace}/resourcequotas/{name}	2002
14.15.2.7. /api/v1/namespaces/{namespace}/resourcequotas/{name}/status	2006
14.16. SECRET [V1]	2009
14.16.1. Specification	2009
14.16.2. API endpoints	2010
14.16.2.1. /api/v1/secrets	2011
14.16.2.2. /api/v1/watch/secrets	2015
14.16.2.3. /api/v1/namespaces/{namespace}/secrets	2018
14.16.2.4. /api/v1/watch/namespaces/{namespace}/secrets	2027
14.16.2.5. /api/v1/namespaces/{namespace}/secrets/{name}	2031
14.16.2.6. /api/v1/watch/namespaces/{namespace}/secrets/{name}	2035
14.17. SERVICE [V1]	2039
14.17.1. Specification	2039
14.17.1.1. .spec	2040
14.17.1.2. .spec.ports	2052
14.17.1.3. .spec.ports[]	2052
14.17.1.4. .spec.sessionAffinityConfig	2054
14.17.1.5. .spec.sessionAffinityConfig.clientIP	2054
14.17.1.6. .status	2054
14.17.1.7. .status.loadBalancer	2054
14.17.1.8. .status.loadBalancer.ingress	2055
14.17.1.9. .status.loadBalancer.ingress[]	2055
14.17.1.10. .status.loadBalancer.ingress[].ports	2056
14.17.1.11. .status.loadBalancer.ingress[].ports[]	2056
14.17.2. API endpoints	2057
14.17.2.1. /api/v1/services	2058
14.17.2.2. /api/v1/watch/services	2061
14.17.2.3. /api/v1/namespaces/{namespace}/services	2064
14.17.2.4. /api/v1/watch/namespaces/{namespace}/services	2073
14.17.2.5. /api/v1/namespaces/{namespace}/services/{name}	2077
14.17.2.6. /api/v1/watch/namespaces/{namespace}/services/{name}	2081
14.17.2.7. /api/v1/namespaces/{namespace}/services/{name}/status	2085
14.18. SERVICEACCOUNT [V1]	2088
14.18.1. Specification	2088
14.18.1.1. .imagePullSecrets	2090
14.18.1.2. .imagePullSecrets[]	2090

14.18.1.3. .secrets	2091
14.18.1.4. .secrets[]	2091
14.18.2. API endpoints	2093
14.18.2.1. /api/v1/serviceaccounts	2093
14.18.2.2. /api/v1/watch/serviceaccounts	2097
14.18.2.3. /api/v1/namespaces/{namespace}/serviceaccounts	2100
14.18.2.4. /api/v1/watch/namespaces/{namespace}/serviceaccounts	2109
14.18.2.5. /api/v1/namespaces/{namespace}/serviceaccounts/{name}	2113
14.18.2.6. /api/v1/watch/namespaces/{namespace}/serviceaccounts/{name}	2117
<b>CHAPTER 15. DISCOVERY APIS</b>	<b>2122</b>
15.1. DISCOVERY APIS	2122
15.1.1. EndpointSlice [discovery.k8s.io/v1]	2122
15.2. ENDOPOINTS LICE [DISCOVERY.K8S.IO/V1]	2122
15.2.1. Specification	2122
15.2.1.1. .endpoints	2124
15.2.1.2. .endpoints[]	2124
15.2.1.3. .endpoints[].conditions	2125
15.2.1.4. .endpoints[].hints	2126
15.2.1.5. .endpoints[].hints.forZones	2127
15.2.1.6. .endpoints[].hints.forZones[]	2127
15.2.1.7. .ports	2127
15.2.1.8. .ports[]	2127
15.2.2. API endpoints	2129
15.2.2.1. /apis/discovery.k8s.io/v1/endpointlices	2130
15.2.2.2. /apis/discovery.k8s.io/v1/watch/endpointlices	2133
15.2.2.3. /apis/discovery.k8s.io/v1/namespaces/{namespace}/endpointlices	2136
15.2.2.4. /apis/discovery.k8s.io/v1/watch/namespaces/{namespace}/endpointlices	2145
15.2.2.5. /apis/discovery.k8s.io/v1/namespaces/{namespace}/endpointlices/{name}	2149
15.2.2.6. /apis/discovery.k8s.io/v1/watch/namespaces/{namespace}/endpointlices/{name}	2153
<b>CHAPTER 16. EVENTS APIS</b>	<b>2158</b>
16.1. EVENTS APIS	2158
16.1.1. Event [events.k8s.io/v1]	2158
16.2. EVENT [EVENTS.K8S.IO/V1]	2158
16.2.1. Specification	2158
16.2.1.1. .series	2161
16.2.2. API endpoints	2161
16.2.2.1. /apis/events.k8s.io/v1/events	2162
16.2.2.2. /apis/events.k8s.io/v1/watch/events	2166
16.2.2.3. /apis/events.k8s.io/v1/namespaces/{namespace}/events	2169
16.2.2.4. /apis/events.k8s.io/v1/watch/namespaces/{namespace}/events	2178
16.2.2.5. /apis/events.k8s.io/v1/namespaces/{namespace}/events/{name}	2182
16.2.2.6. /apis/events.k8s.io/v1/watch/namespaces/{namespace}/events/{name}	2186
<b>CHAPTER 17. FLOW CONTROL APIS</b>	<b>2191</b>
17.1. FLOW CONTROL APIS	2191
17.1.1. FlowSchema [flowcontrol.apiserver.k8s.io/v1beta3]	2191
17.1.2. PriorityLevelConfiguration [flowcontrol.apiserver.k8s.io/v1beta3]	2191
17.2. FLOWSHEMA [FLOWCONTROL.APISERVER.K8S.IO/V1BETA3]	2191
17.2.1. Specification	2191
17.2.1.1. .spec	2192
17.2.1.2. .spec.distinguisherMethod	2193
17.2.1.3. .spec.priorityLevelConfiguration	2194

17.2.1.4. .spec.rules	2194
17.2.1.5. .spec.rules[]	2194
17.2.1.6. .spec.rules[].nonResourceRules	2196
17.2.1.7. .spec.rules[].nonResourceRules[]	2196
17.2.1.8. .spec.rules[].resourceRules	2197
17.2.1.9. .spec.rules[].resourceRules[]	2197
17.2.1.10. .spec.rules[].subjects	2199
17.2.1.11. .spec.rules[].subjects[]	2199
17.2.1.12. .spec.rules[].subjects[].group	2200
17.2.1.13. .spec.rules[].subjects[].serviceAccount	2200
17.2.1.14. .spec.rules[].subjects[].user	2201
17.2.1.15. .status	2201
17.2.1.16. .status.conditions	2201
17.2.1.17. .status.conditions[]	2202
17.2.2. API endpoints	2202
17.2.2.1. /apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas	2203
17.2.2.2. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/flowschemas	2212
17.2.2.3. /apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas/{name}	2215
17.2.2.4. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/flowschemas/{name}	2219
17.2.2.5. /apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas/{name}/status	2223
17.3. PRIORITYLEVELCONFIGURATION [FLOWCONTROL.APISERVER.K8S.IO/V1BETA3]	2226
17.3.1. Specification	2226
17.3.1.1. .spec	2227
17.3.1.2. .spec.limited	2228
17.3.1.3. .spec.limited.limitResponse	2230
17.3.1.4. .spec.limited.limitResponse.queuing	2231
17.3.1.5. .status	2232
17.3.1.6. .status.conditions	2232
17.3.1.7. .status.conditions[]	2232
17.3.2. API endpoints	2233
17.3.2.1. /apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations	2234
17.3.2.2. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/prioritylevelconfigurations	2243
17.3.2.3. /apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations/{name}	2246
17.3.2.4. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/prioritylevelconfigurations/{name}	2250
17.3.2.5. /apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations/{name}/status	2254
<b>CHAPTER 18. NETWORKING APIS</b> .....	<b>2258</b>
18.1. NETWORKING APIS	2258
18.1.1. Ingress [networking.k8s.io/v1]	2258
18.1.2. IngressClass [networking.k8s.io/v1]	2258
18.1.3. NetworkPolicy [networking.k8s.io/v1]	2258
18.2. INGRESS [NETWORKING.K8S.IO/V1]	2258
18.2.1. Specification	2258
18.2.1.1. .spec	2259
18.2.1.2. .spec.defaultBackend	2261
18.2.1.3. .spec.defaultBackend.service	2261
18.2.1.4. .spec.defaultBackend.service.port	2262
18.2.1.5. .spec.rules	2262
18.2.1.6. .spec.rules[]	2262
18.2.1.7. .spec.rules[].http	2264
18.2.1.8. .spec.rules[].http.paths	2264
18.2.1.9. .spec.rules[].http.paths[]	2264
18.2.1.10. .spec.rules[].http.paths[].backend	2266

18.2.1.11. .spec.rules[].http.paths[].backend.service	2267
18.2.1.12. .spec.rules[].http.paths[].backend.service.port	2267
18.2.1.13. .spec.tls	2268
18.2.1.14. .spec.tls[]	2268
18.2.1.15. .status	2269
18.2.1.16. .status.loadBalancer	2269
18.2.1.17. .status.loadBalancer.ingress	2269
18.2.1.18. .status.loadBalancer.ingress[]	2269
18.2.1.19. .status.loadBalancer.ingress[].ports	2270
18.2.1.20. .status.loadBalancer.ingress[].ports[]	2270
18.2.2. API endpoints	2271
18.2.2.1. /apis/networking.k8s.io/v1/ingresses	2272
18.2.2.2. /apis/networking.k8s.io/v1/watch/ingresses	2276
18.2.2.3. /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses	2279
18.2.2.4. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/ingresses	2288
18.2.2.5. /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}	2292
18.2.2.6. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/ingresses/{name}	2296
18.2.2.7. /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}/status	2300
18.3. INGRESSCLASS [NETWORKING.K8S.IO/V1]	2303
18.3.1. Specification	2303
18.3.1.1. .spec	2304
18.3.1.2. .spec.parameters	2305
18.3.2. API endpoints	2306
18.3.2.1. /apis/networking.k8s.io/v1/ingressclasses	2306
18.3.2.2. /apis/networking.k8s.io/v1/watch/ingressclasses	2315
18.3.2.3. /apis/networking.k8s.io/v1/ingressclasses/{name}	2318
18.3.2.4. /apis/networking.k8s.io/v1/watch/ingressclasses/{name}	2322
18.4. NETWORKPOLICY [NETWORKING.K8S.IO/V1]	2326
18.4.1. Specification	2326
18.4.1.1. .spec	2327
18.4.1.2. .spec.egress	2329
18.4.1.3. .spec.egress[]	2329
18.4.1.4. .spec.egress[].ports	2330
18.4.1.5. .spec.egress[].ports[]	2330
18.4.1.6. .spec.egress[].to	2331
18.4.1.7. .spec.egress[].to[]	2331
18.4.1.8. .spec.egress[].to[].ipBlock	2332
18.4.1.9. .spec.ingress	2333
18.4.1.10. .spec.ingress[]	2333
18.4.1.11. .spec.ingress[].from	2334
18.4.1.12. .spec.ingress[].from[]	2334
18.4.1.13. .spec.ingress[].from[].ipBlock	2335
18.4.1.14. .spec.ingress[].ports	2336
18.4.1.15. .spec.ingress[].ports[]	2336
18.4.1.16. .status	2337
18.4.2. API endpoints	2337
18.4.2.1. /apis/networking.k8s.io/v1/networkpolicies	2338
18.4.2.2. /apis/networking.k8s.io/v1/watch/networkpolicies	2342
18.4.2.3. /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies	2345
18.4.2.4. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/networkpolicies	2354
18.4.2.5. /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}	2358
18.4.2.6. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/networkpolicies/{name}	2362
18.4.2.7. /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}/status	2366

<b>CHAPTER 19. NODE APIS</b> .....	<b>2370</b>
19.1. NODE APIS	2370
19.1.1. RuntimeClass [node.k8s.io/v1]	2370
19.2. RUNTIMECLASS [NODE.K8S.IO/V1]	2370
19.2.1. Specification	2370
19.2.1.1. .overhead	2371
19.2.1.2. .scheduling	2372
19.2.2. API endpoints	2372
19.2.2.1. /apis/node.k8s.io/v1/runtimeclasses	2373
19.2.2.2. /apis/node.k8s.io/v1/watch/runtimeclasses	2382
19.2.2.3. /apis/node.k8s.io/v1/runtimeclasses/{name}	2385
19.2.2.4. /apis/node.k8s.io/v1/watch/runtimeclasses/{name}	2389
<b>CHAPTER 20. POLICY APIS</b> .....	<b>2394</b>
20.1. POLICY APIS	2394
20.1.1. Eviction [policy/v1]	2394
20.1.2. PodDisruptionBudget [policy/v1]	2394
20.2. EVICTION [POLICY/V1]	2394
20.2.1. Specification	2394
20.2.2. API endpoints	2395
20.2.2.1. /api/v1/namespaces/{namespace}/pods/{name}/eviction	2395
20.3. PODDISRUPTIONBUDGET [POLICY/V1]	2397
20.3.1. Specification	2397
20.3.1.1. .spec	2398
20.3.1.2. .status	2400
20.3.2. API endpoints	2402
20.3.2.1. /apis/policy/v1/poddisruptionbudgets	2403
20.3.2.2. /apis/policy/v1/watch/poddisruptionbudgets	2407
20.3.2.3. /apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets	2410
20.3.2.4. /apis/policy/v1/watch/namespaces/{namespace}/poddisruptionbudgets	2419
20.3.2.5. /apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets/{name}	2423
20.3.2.6. /apis/policy/v1/watch/namespaces/{namespace}/poddisruptionbudgets/{name}	2427
20.3.2.7. /apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets/{name}/status	2431
<b>CHAPTER 21. RBAC APIS</b> .....	<b>2435</b>
21.1. RBAC APIS	2435
21.1.1. ClusterRole [rbac.authorization.k8s.io/v1]	2435
21.1.2. ClusterRoleBinding [rbac.authorization.k8s.io/v1]	2435
21.1.3. Role [rbac.authorization.k8s.io/v1]	2435
21.1.4. RoleBinding [rbac.authorization.k8s.io/v1]	2435
21.2. CLUSTERROLE [RBAC.AUTHORIZATION.K8S.IO/V1]	2435
21.2.1. Specification	2435
21.2.1.1. .aggregationRule	2436
21.2.1.2. .rules	2437
21.2.1.3. .rules[]	2437
21.2.2. API endpoints	2438
21.2.2.1. /apis/rbac.authorization.k8s.io/v1/clusterroles	2439
21.2.2.2. /apis/rbac.authorization.k8s.io/v1/watch/clusterroles	2447
21.2.2.3. /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}	2450
21.2.2.4. /apis/rbac.authorization.k8s.io/v1/watch/clusterroles/{name}	2454
21.3. CLUSTERROLEBINDING [RBAC.AUTHORIZATION.K8S.IO/V1]	2458
21.3.1. Specification	2458
21.3.1.1. .roleRef	2459

21.3.1.2. .subjects	2459
21.3.1.3. .subjects[]	2459
21.3.2. API endpoints	2460
21.3.2.1. /apis/rbac.authorization.k8s.io/v1/clusterrolebindings	2461
21.3.2.2. /apis/rbac.authorization.k8s.io/v1/watch/clusterrolebindings	2470
21.3.2.3. /apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}	2473
21.3.2.4. /apis/rbac.authorization.k8s.io/v1/watch/clusterrolebindings/{name}	2477
21.4. ROLE [RBAC.AUTHORIZATION.K8S.IO/V1]	2481
21.4.1. Specification	2481
21.4.1.1. .rules	2482
21.4.1.2. .rules[]	2482
21.4.2. API endpoints	2483
21.4.2.1. /apis/rbac.authorization.k8s.io/v1/roles	2484
21.4.2.2. /apis/rbac.authorization.k8s.io/v1/watch/roles	2488
21.4.2.3. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles	2491
21.4.2.4. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/roles	2500
21.4.2.5. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}	2504
21.4.2.6. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/roles/{name}	2508
21.5. ROLEBINDING [RBAC.AUTHORIZATION.K8S.IO/V1]	2512
21.5.1. Specification	2512
21.5.1.1. .roleRef	2513
21.5.1.2. .subjects	2513
21.5.1.3. .subjects[]	2514
21.5.2. API endpoints	2514
21.5.2.1. /apis/rbac.authorization.k8s.io/v1/rolebindings	2515
21.5.2.2. /apis/rbac.authorization.k8s.io/v1/watch/rolebindings	2519
21.5.2.3. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings	2522
21.5.2.4. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/rolebindings	2531
21.5.2.5. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}	2535
21.5.2.6. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/rolebindings/{name}	2539
<b>CHAPTER 22. NETWORK APIS</b> .....	<b>2544</b>
22.1. NETWORK APIS	2544
22.1.1. Route [route.openshift.io/v1]	2544
22.2. ROUTE [ROUTE.OPENSIFT.IO/V1]	2544
22.2.1. Specification	2545
22.2.1.1. .spec	2545
22.2.1.2. .spec.alternateBackends	2547
22.2.1.3. .spec.alternateBackends[]	2548
22.2.1.4. .spec.port	2548
22.2.1.5. .spec.tls	2549
22.2.1.6. .spec.to	2550
22.2.1.7. .status	2551
22.2.1.8. .status.ingress	2551
22.2.1.9. .status.ingress[]	2551
22.2.1.10. .status.ingress[].conditions	2552
22.2.1.11. .status.ingress[].conditions[]	2552
22.2.2. API endpoints	2553
22.2.2.1. /apis/route.openshift.io/v1/routes	2554
22.2.2.2. /apis/route.openshift.io/v1/namespaces/{namespace}/routes	2557
22.2.2.3. /apis/route.openshift.io/v1/namespaces/{namespace}/routes/{name}	2565
22.2.2.4. /apis/route.openshift.io/v1/namespaces/{namespace}/routes/{name}/status	2570



<b>CHAPTER 23. SCHEDULING APIS</b> .....	<b>2574</b>
23.1. SCHEDULING APIS	2574
23.1.1. PriorityClass [scheduling.k8s.io/v1]	2574
23.2. PRIORITYCLASS [SCHEDULING.K8S.IO/V1]	2574
23.2.1. Specification	2574
23.2.2. API endpoints	2576
23.2.2.1. /apis/scheduling.k8s.io/v1/priorityclasses	2576
23.2.2.2. /apis/scheduling.k8s.io/v1/watch/priorityclasses	2585
23.2.2.3. /apis/scheduling.k8s.io/v1/priorityclasses/{name}	2588
23.2.2.4. /apis/scheduling.k8s.io/v1/watch/priorityclasses/{name}	2592
<b>CHAPTER 24. SECURITY APIS</b> .....	<b>2597</b>
24.1. SECURITY APIS	2597
24.1.1. SecurityContextConstraints [security.openshift.io/v1]	2597
24.2. SECURITYCONTEXTCONSTRAINTS [SECURITY.OPENSIFT.IO/V1]	2597
24.2.1. Specification	2597
24.2.2. API endpoints	2602
24.2.2.1. /apis/security.openshift.io/v1/securitycontextconstraints	2602
24.2.2.2. /apis/security.openshift.io/v1/securitycontextconstraints/{name}	2610
<b>CHAPTER 25. SECURITY-INTERNAL APIS</b> .....	<b>2615</b>
25.1. SECURITY INTERNAL APIS	2615
25.1.1. RangeAllocation [security.internal.openshift.io/v1]	2615
25.2. RANGEALLOCATION [SECURITY.INTERNAL.OPENSIFT.IO/V1]	2615
25.2.1. Specification	2615
25.2.2. API endpoints	2616
25.2.2.1. /apis/security.internal.openshift.io/v1/rangeallocations	2616
25.2.2.2. /apis/security.internal.openshift.io/v1/rangeallocations/{name}	2624
<b>CHAPTER 26. SNAPSHOT APIS</b> .....	<b>2629</b>
26.1. SNAPSHOT APIS	2629
26.1.1. VolumeSnapshot [snapshot.storage.k8s.io/v1]	2629
26.1.2. VolumeSnapshotClass [snapshot.storage.k8s.io/v1]	2629
26.1.3. VolumeSnapshotContent [snapshot.storage.k8s.io/v1]	2629
26.2. VOLUMESNAPSHOT [SNAPSHOT.STORAGE.K8S.IO/V1]	2629
26.2.1. Specification	2629
26.2.1.1. .spec	2630
26.2.1.2. .spec.source	2631
26.2.1.3. .status	2632
26.2.1.4. .status.error	2635
26.2.2. API endpoints	2635
26.2.2.1. /apis/snapshot.storage.k8s.io/v1/volumesnapshots	2636
26.2.2.2. /apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots	2640
26.2.2.3. /apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots/{name}	2648
26.2.2.4. /apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots/{name}/status	2653
26.3. VOLUMESNAPSHOTCLASS [SNAPSHOT.STORAGE.K8S.IO/V1]	2656
26.3.1. Specification	2656
26.3.2. API endpoints	2657
26.3.2.1. /apis/snapshot.storage.k8s.io/v1/volumesnapshotclasses	2658
26.3.2.2. /apis/snapshot.storage.k8s.io/v1/volumesnapshotclasses/{name}	2666
26.4. VOLUMESNAPSHOTCONTENT [SNAPSHOT.STORAGE.K8S.IO/V1]	2671
26.4.1. Specification	2671
26.4.1.1. .spec	2672

26.4.1.2. .spec.source	2674
26.4.1.3. .spec.volumeSnapshotRef	2675
26.4.1.4. .status	2677
26.4.1.5. .status.error	2678
26.4.2. API endpoints	2679
26.4.2.1. /apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents	2679
26.4.2.2. /apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents/{name}	2687
26.4.2.3. /apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents/{name}/status	2692
<b>CHAPTER 27. STORAGE APIS</b>	<b>2696</b>
27.1. STORAGE APIS	2696
27.1.1. CSIDriver [storage.k8s.io/v1]	2696
27.1.2. CSINode [storage.k8s.io/v1]	2696
27.1.3. CSIStorageCapacity [storage.k8s.io/v1]	2696
27.1.4. StorageClass [storage.k8s.io/v1]	2697
27.1.5. VolumeAttachment [storage.k8s.io/v1]	2697
27.2. CSIDRIVER [STORAGE.K8S.IO/V1]	2697
27.2.1. Specification	2697
27.2.1.1. .spec	2698
27.2.1.2. .spec.tokenRequests	2703
27.2.1.3. .spec.tokenRequests[]	2704
27.2.2. API endpoints	2704
27.2.2.1. /apis/storage.k8s.io/v1/csidders	2705
27.2.2.2. /apis/storage.k8s.io/v1/watch/csidders	2713
27.2.2.3. /apis/storage.k8s.io/v1/csidders/{name}	2716
27.2.2.4. /apis/storage.k8s.io/v1/watch/csidders/{name}	2720
27.3. CSINODE [STORAGE.K8S.IO/V1]	2724
27.3.1. Specification	2724
27.3.1.1. .spec	2725
27.3.1.2. .spec.drivers	2725
27.3.1.3. .spec.drivers[]	2725
27.3.1.4. .spec.drivers[].allocatable	2727
27.3.2. API endpoints	2727
27.3.2.1. /apis/storage.k8s.io/v1/csinodes	2728
27.3.2.2. /apis/storage.k8s.io/v1/watch/csinodes	2737
27.3.2.3. /apis/storage.k8s.io/v1/csinodes/{name}	2740
27.3.2.4. /apis/storage.k8s.io/v1/watch/csinodes/{name}	2744
27.4. CSISTORAGECAPACITY [STORAGE.K8S.IO/V1]	2748
27.4.1. Specification	2748
27.4.2. API endpoints	2750
27.4.2.1. /apis/storage.k8s.io/v1/csistoragecapacities	2751
27.4.2.2. /apis/storage.k8s.io/v1/watch/csistoragecapacities	2755
27.4.2.3. /apis/storage.k8s.io/v1/namespaces/{namespace}/csistoragecapacities	2758
27.4.2.4. /apis/storage.k8s.io/v1/watch/namespaces/{namespace}/csistoragecapacities	2767
27.4.2.5. /apis/storage.k8s.io/v1/namespaces/{namespace}/csistoragecapacities/{name}	2771
27.4.2.6. /apis/storage.k8s.io/v1/watch/namespaces/{namespace}/csistoragecapacities/{name}	2775
27.5. STORAGECLASS [STORAGE.K8S.IO/V1]	2779
27.5.1. Specification	2779
27.5.2. API endpoints	2781
27.5.2.1. /apis/storage.k8s.io/v1/storageclasses	2782
27.5.2.2. /apis/storage.k8s.io/v1/watch/storageclasses	2791
27.5.2.3. /apis/storage.k8s.io/v1/storageclasses/{name}	2794
27.5.2.4. /apis/storage.k8s.io/v1/watch/storageclasses/{name}	2798

27.6. VOLUMEATTACHMENT [STORAGE.K8S.IO/V1]	2802
27.6.1. Specification	2802
27.6.1.1. .spec	2803
27.6.1.2. .spec.source	2804
27.6.1.3. .status	2804
27.6.1.4. .status.attachError	2805
27.6.1.5. .status.detachError	2805
27.6.2. API endpoints	2806
27.6.2.1. /apis/storage.k8s.io/v1/volumeattachments	2807
27.6.2.2. /apis/storage.k8s.io/v1/watch/volumeattachments	2815
27.6.2.3. /apis/storage.k8s.io/v1/volumeattachments/{name}	2818
27.6.2.4. /apis/storage.k8s.io/v1/watch/volumeattachments/{name}	2822
27.6.2.5. /apis/storage.k8s.io/v1/volumeattachments/{name}/status	2826
<b>CHAPTER 28. STORAGE VERSION MIGRATION APIS</b> .....	<b>2830</b>
28.1. STORAGE VERSION MIGRATION APIS	2830
28.1.1. StorageVersionMigration [migration.k8s.io/v1alpha1]	2830
28.2. STORAGEVERSIONMIGRATION [MIGRATION.K8S.IO/V1ALPHA1]	2830
28.2.1. Specification	2830
28.2.1.1. .spec	2831
28.2.1.2. .spec.resource	2831
28.2.1.3. .status	2832
28.2.1.4. .status.conditions	2832
28.2.1.5. .status.conditions[]	2832
28.2.2. API endpoints	2833
28.2.2.1. /apis/migration.k8s.io/v1alpha1/storageversionmigrations	2833
28.2.2.2. /apis/migration.k8s.io/v1alpha1/storageversionmigrations/{name}	2841
28.2.2.3. /apis/migration.k8s.io/v1alpha1/storageversionmigrations/{name}/status	2846
<b>CHAPTER 29. TOPOLVM APIS</b> .....	<b>2850</b>
29.1. TOPOLVM APIS	2850
29.1.1. LogicalVolume [topolvm.io/v1]	2850
29.2. LOGICALVOLUME [TOPOLVM.IO/V1]	2850
29.2.1. Specification	2850
29.2.1.1. .spec	2851
29.2.1.2. .status	2852
29.2.2. API endpoints	2852
29.2.2.1. /apis/topolvm.io/v1/logicalvolumes	2853
29.2.2.2. /apis/topolvm.io/v1/logicalvolumes/{name}	2861
29.2.2.3. /apis/topolvm.io/v1/logicalvolumes/{name}/status	2866
<b>CHAPTER 30. WEBHOOK APIS</b> .....	<b>2870</b>
30.1. WEBHOOK APIS	2870
30.1.1. MutatingWebhookConfiguration [admissionregistration.k8s.io/v1]	2870
30.1.2. ValidatingWebhookConfiguration [admissionregistration.k8s.io/v1]	2870
30.2. MUTATINGWEBHOOKCONFIGURATION [ADMISSIONREGISTRATION.K8S.IO/V1]	2870
30.2.1. Specification	2870
30.2.1.1. .webhooks	2871
30.2.1.2. .webhooks[]	2871
30.2.1.3. .webhooks[].clientConfig	2879
30.2.1.4. .webhooks[].clientConfig.service	2881
30.2.1.5. .webhooks[].matchConditions	2882
30.2.1.6. .webhooks[].matchConditions[]	2882
30.2.1.7. .webhooks[].rules	2884

30.2.1.8. .webhooks[].rules[]	2884
30.2.2. API endpoints	2886
30.2.2.1. /apis/admissionregistration.k8s.io/v1/mutatingwebhookconfigurations	2886
30.2.2.2. /apis/admissionregistration.k8s.io/v1/watch/mutatingwebhookconfigurations	2895
30.2.2.3. /apis/admissionregistration.k8s.io/v1/mutatingwebhookconfigurations/{name}	2898
30.2.2.4. /apis/admissionregistration.k8s.io/v1/watch/mutatingwebhookconfigurations/{name}	2902
30.3. VALIDATINGWEBHOOKCONFIGURATION [ADMISSIONREGISTRATION.K8S.IO/V1]	2906
30.3.1. Specification	2906
30.3.1.1. .webhooks	2907
30.3.1.2. .webhooks[]	2907
30.3.1.3. .webhooks[].clientConfig	2913
30.3.1.4. .webhooks[].clientConfig.service	2915
30.3.1.5. .webhooks[].matchConditions	2916
30.3.1.6. .webhooks[].matchConditions[]	2916
30.3.1.7. .webhooks[].rules	2918
30.3.1.8. .webhooks[].rules[]	2918
30.3.2. API endpoints	2920
30.3.2.1. /apis/admissionregistration.k8s.io/v1/validatingwebhookconfigurations	2920
30.3.2.2. /apis/admissionregistration.k8s.io/v1/watch/validatingwebhookconfigurations	2929
30.3.2.3. /apis/admissionregistration.k8s.io/v1/validatingwebhookconfigurations/{name}	2932
30.3.2.4. /apis/admissionregistration.k8s.io/v1/watch/validatingwebhookconfigurations/{name}	2936



# CHAPTER 1. UNDERSTANDING API TIERS



## IMPORTANT

This guidance does not cover layered Red Hat build of MicroShift offerings.

Red Hat requests that application developers validate that any behavior they depend on is explicitly defined in the formal API documentation to prevent introducing dependencies on unspecified implementation-specific behavior or dependencies on bugs in a particular implementation of an API. For example, new releases of an ingress router may not be compatible with older releases if an application uses an undocumented API or relies on undefined behavior.

## 1.1. API TIERS

All commercially supported APIs, components, and features are associated under one of the following support levels:

### API tier 1

APIs and application operating environments (AOEs) are stable within a major release. They may be deprecated within a major release, but they will not be removed until a subsequent major release.

### API tier 2

APIs and AOEs are stable within a major release for a minimum of 9 months or 3 minor releases from the announcement of deprecation, whichever is longer.

### API tier 3

This level applies to languages, tools, applications, and optional Operators included with Red Hat build of MicroShift through Operator Hub. Each component will specify a lifetime during which the API and AOE will be supported. Newer versions of language runtime specific components will attempt to be as API and AOE compatible from minor version to minor version as possible. Minor version to minor version compatibility is not guaranteed, however.

Components and developer tools that receive continuous updates through the Operator Hub, referred to as Operators and operands, should be considered API tier 3. Developers should use caution and understand how these components may change with each minor release. Users are encouraged to consult the compatibility guidelines documented by the component.

### API tier 4

No compatibility is provided. API and AOE can change at any point. These capabilities should not be used by applications needing long-term support.

It is common practice for Operators to use custom resource definitions (CRDs) internally to accomplish a task. These objects are not meant for use by actors external to the Operator and are intended to be hidden. If any CRD is not meant for use by actors external to the Operator, the **operators.operatorframework.io/internal-objects** annotation in the Operators **ClusterServiceVersion** (CSV) should be specified to signal that the corresponding resource is internal use only and the CRD may be explicitly labeled as tier 4.

## 1.2. MAPPING API TIERS TO API GROUPS

For each API tier defined by Red Hat, we provide a mapping table for specific API groups where the upstream communities are committed to maintain forward compatibility. Any API group that does not specify an explicit compatibility level and is not specifically discussed below is assigned API tier 3 by default except for **v1alpha1** APIs which are assigned tier 4 by default.

### 1.2.1. Support for Kubernetes API groups

API groups that end with the suffix **\*.k8s.io** or have the form **version.<name>** with no suffix are governed by the Kubernetes deprecation policy and follow a general mapping between API version exposed and corresponding support tier unless otherwise specified.

API version example	API tier
<b>v1</b>	Tier 1
<b>v1beta1</b>	Tier 2
<b>v1alpha1</b>	Tier 4

### 1.2.2. Support for OpenShift API groups

API groups that end with the suffix **\*.openshift.io** are governed by the Red Hat build of MicroShift deprecation policy and follow a general mapping between API version exposed and corresponding compatibility level unless otherwise specified.

API version example	API tier
<b>route.openshift.io/v1</b>	Tier 1
<b>security.openshift.io/v1</b>	Tier 1 except for <b>RangeAllocation</b> (tier 4) and <b>*Reviews</b> (tier 2)

## 1.3. API DEPRECATION POLICY

Red Hat build of MicroShift is composed of many components sourced from many upstream communities. It is anticipated that the set of components, the associated API interfaces, and correlated features will evolve over time and might require formal deprecation in order to remove the capability.

### 1.3.1. Deprecating parts of the API

Red Hat build of MicroShift is a distributed system where multiple components interact with a shared state managed by the cluster control plane through a set of structured APIs. Per Kubernetes conventions, each API presented by Red Hat build of MicroShift is associated with a group identifier and each API group is independently versioned. Each API group is managed in a distinct upstream community including Kubernetes, Metal3, Multus, Operator Framework, Open Cluster Management, OpenShift itself, and more.

While each upstream community might define their own unique deprecation policy for a given API group and version, Red Hat normalizes the community specific policy to one of the compatibility levels defined prior based on our integration in and awareness of each upstream community to simplify end-user consumption and support.

The deprecation policy and schedule for APIs vary by compatibility level.

The deprecation policy covers all elements of the API including:

- REST resources, also known as API objects
- Fields of REST resources
- Annotations on REST resources, excluding version-specific qualifiers
- Enumerated or constant values

Other than the most recent API version in each group, older API versions must be supported after their announced deprecation for a duration of no less than:

API tier	Duration
Tier 1	Stable within a major release. They may be deprecated within a major release, but they will not be removed until a subsequent major release.
Tier 2	9 months or 3 releases from the announcement of deprecation, whichever is longer.
Tier 3	See the component-specific schedule.
Tier 4	None. No compatibility is guaranteed.

The following rules apply to all tier 1 APIs:

- API elements can only be removed by incrementing the version of the group.
- API objects must be able to round-trip between API versions without information loss, with the exception of whole REST resources that do not exist in some versions. In cases where equivalent fields do not exist between versions, data will be preserved in the form of annotations during conversion.
- API versions in a given group can not deprecate until a new API version at least as stable is released, except in cases where the entire API object is being removed.

### 1.3.2. Deprecating CLI elements

Client-facing CLI commands are not versioned in the same way as the API, but are user-facing component systems. The two major ways a user interacts with a CLI are through a command or flag, which is referred to in this context as CLI elements.

All CLI elements default to API tier 1 unless otherwise noted or the CLI depends on a lower tier API.

	Element	API tier
Generally available (GA)	Flags and commands	Tier 1
Technology Preview	Flags and commands	Tier 3
Developer Preview	Flags and commands	Tier 4



### 1.3.3. Deprecating an entire component

The duration and schedule for deprecating an entire component maps directly to the duration associated with the highest API tier of an API exposed by that component. For example, a component that surfaced APIs with tier 1 and 2 could not be removed until the tier 1 deprecation schedule was met.

API tier	Duration
Tier 1	Stable within a major release. They may be deprecated within a major release, but they will not be removed until a subsequent major release.
Tier 2	9 months or 3 releases from the announcement of deprecation, whichever is longer.
Tier 3	See the component-specific schedule.
Tier 4	None. No compatibility is guaranteed.

## CHAPTER 2. UNDERSTANDING API COMPATIBILITY GUIDELINES

Follow the compatibility guidelines to understand the APIs enabled for Red Hat build of MicroShift.



### IMPORTANT

This guidance does not cover layered Red Hat build of MicroShift offerings.

### 2.1. API COMPATIBILITY GUIDELINES

Red Hat recommends that application developers adopt the following principles in order to improve compatibility with Red Hat build of MicroShift:

- Use APIs and components with support tiers that match the application's need.
- Build applications using the published client libraries where possible.
- Applications are only guaranteed to run correctly if they execute in an environment that is as new as the environment it was built to execute against. An application that was built for Red Hat build of MicroShift 4.14 is not guaranteed to function properly on Red Hat build of MicroShift 4.13.
- Do not design applications that rely on configuration files provided by system packages or other components. These files can change between versions unless the upstream community is explicitly committed to preserving them. Where appropriate, depend on any Red Hat provided interface abstraction over those configuration files in order to maintain forward compatibility. Direct file system modification of configuration files is discouraged, and users are strongly encouraged to integrate with an Operator provided API where available to avoid dual-writer conflicts.
- Do not depend on API fields prefixed with **unsupported<FieldName>** or annotations that are not explicitly mentioned in product documentation.
- Do not depend on components with shorter compatibility guarantees than your application.
- Do not perform direct storage operations on the etcd server. All etcd access must be performed via the api-server or through documented backup and restore procedures.

Red Hat recommends that application developers follow the [compatibility guidelines](#) defined by Red Hat Enterprise Linux (RHEL). Red Hat build of MicroShift strongly recommends the following guidelines when building an application or hosting an application on the platform:

- Do not depend on a specific Linux kernel or Red Hat build of MicroShift version.
- Avoid reading from **proc**, **sys**, and **debug** file systems, or any other pseudo file system.
- Avoid using **ioctl**s to directly interact with hardware.
- Avoid direct interaction with **cgroups** in order to not conflict with Red Hat build of MicroShift host-agents that provide the container execution environment.

**NOTE**

During the lifecycle of a release, Red Hat makes commercially reasonable efforts to maintain API and application operating environment (AOE) compatibility across all minor releases and z-stream releases. If necessary, Red Hat might make exceptions to this compatibility goal for critical impact security or other significant issues.

## 2.2. API COMPATIBILITY EXCEPTIONS

The following are exceptions to compatibility in Red Hat build of MicroShift:

### Functional defaults between an upgraded cluster and a new installation

No assurances are made at this time that a new installation of a product minor release will have the same functional defaults as a version of the product that was installed with a prior minor release and upgraded to the equivalent version. For example, future versions of the product may provision cloud infrastructure with different defaults than prior minor versions. In addition, different default security choices may be made in future versions of the product than those made in past versions of the product. Past versions of the product will forward upgrade, but preserve legacy choices where appropriate specifically to maintain backwards compatibility.

### Usage of API fields that have the prefix "unsupported" or undocumented annotations

Select APIs in the product expose fields with the prefix **unsupported<FieldName>**. No assurances are made at this time that usage of this field is supported across releases or within a release. Product support can request a customer to specify a value in this field when debugging specific problems, but its usage is not supported outside of that interaction. Usage of annotations on objects that are not explicitly documented are not assured support across minor releases.

### API availability per product installation topology

The OpenShift distribution will continue to evolve its supported installation topology, and not all APIs in one install topology will necessarily be included in another. For example, certain topologies may restrict read/write access to particular APIs if they are in conflict with the product installation topology or not include a particular API at all if not pertinent to that topology. APIs that exist in a given topology will be supported in accordance with the compatibility tiers defined above.

## 2.3. API COMPATIBILITY COMMON TERMINOLOGY

### 2.3.1. Application Programming Interface (API)

An API is a public interface implemented by a software program that enables it to interact with other software. In Red Hat build of MicroShift, the API is served from a centralized API server and is used as the hub for all system interaction.

### 2.3.2. Application Operating Environment (AOE)

An AOE is the integrated environment that executes the end-user application program. The AOE is a containerized environment that provides isolation from the host operating system (OS). At a minimum, AOE allows the application to run in an isolated manner from the host OS libraries and binaries, but still share the same OS kernel as all other containers on the host. The AOE is enforced at runtime and it describes the interface between an application and its operating environment. It includes intersection points between the platform, operating system and environment, with the user application including projection of downward API, DNS, resource accounting, device access, platform workload identity, isolation among containers, isolation between containers and host OS.

The AOE does not include components that might vary by installation, such as Container Network

Interface (CNI) plugin selection or extensions to the product such as admission hooks. Components that integrate with the cluster at a level below the container environment might be subjected to additional variation between versions.

### 2.3.3. Compatibility in a virtualized environment

Virtual environments emulate bare-metal environments such that unprivileged applications that run on bare-metal environments will run, unmodified, in corresponding virtual environments. Virtual environments present simplified abstracted views of physical resources, so some differences might exist.

### 2.3.4. Compatibility in a cloud environment

Red Hat build of MicroShift might choose to offer integration points with a hosting cloud environment via cloud provider specific integrations. The compatibility of these integration points are specific to the guarantee provided by the native cloud vendor and its intersection with the Red Hat build of MicroShift compatibility window. Where Red Hat build of MicroShift provides an integration with a cloud environment natively as part of the default installation, Red Hat develops against stable cloud API endpoints to provide commercially reasonable support with forward looking compatibility that includes stable deprecation policies. Example areas of integration between the cloud provider and Red Hat build of MicroShift include, but are not limited to, dynamic volume provisioning, service load balancer integration, pod workload identity, dynamic management of compute, and infrastructure provisioned as part of initial installation.

### 2.3.5. Major, minor, and z-stream releases

A Red Hat major release represents a significant step in the development of a product. Minor releases appear more frequently within the scope of a major release and represent deprecation boundaries that might impact future application compatibility. A z-stream release is an update to a minor release which provides a stream of continuous fixes to an associated minor release. API and AOE compatibility is never broken in a z-stream release except when this policy is explicitly overridden in order to respond to an unforeseen security impact.

For example, in the release 4.13.2:

- 4 is the major release version
- 13 is the minor release version
- 2 is the z-stream release version

### 2.3.6. Extended user support (EUS)

A minor release in an Red Hat build of MicroShift major release that has an extended support window for critical bug fixes. Users are able to migrate between EUS releases by incrementally adopting minor versions between EUS releases. It is important to note that the deprecation policy is defined across minor releases and not EUS releases. As a result, an EUS user might have to respond to a deprecation when migrating to a future EUS while sequentially upgrading through each minor release.

### 2.3.7. Developer Preview

An optional product capability that is not officially supported by Red Hat, but is intended to provide a mechanism to explore early phase technology. By default, Developer Preview functionality is opt-in, and subject to removal at any time. Enabling a Developer Preview feature might render a cluster unsupported dependent upon the scope of the feature.

### 2.3.8. Technology Preview

An optional product capability that provides early access to upcoming product innovations to test functionality and provide feedback during the development process. The feature is not fully supported, might not be functionally complete, and is not intended for production use. Usage of a Technology Preview function requires explicit opt-in. Learn more about the [Technology Preview Features Support Scope](#).

## CHAPTER 3. API INDEX

API	API group
<a href="#">APIService</a>	apiregistration.k8s.io/v1
<a href="#">Binding</a>	v1
<a href="#">CertificateSigningRequest</a>	certificates.k8s.io/v1
<a href="#">ClusterRole</a>	rbac.authorization.k8s.io/v1
<a href="#">ClusterRoleBinding</a>	rbac.authorization.k8s.io/v1
<a href="#">ComponentStatus</a>	v1
<a href="#">ConfigMap</a>	v1
<a href="#">ControllerRevision</a>	apps/v1
<a href="#">CronJob</a>	batch/v1
<a href="#">CSIDriver</a>	storage.k8s.io/v1
<a href="#">CSINode</a>	storage.k8s.io/v1
<a href="#">CSIStorageCapacity</a>	storage.k8s.io/v1
<a href="#">CustomResourceDefinition</a>	apiextensions.k8s.io/v1
<a href="#">DaemonSet</a>	apps/v1
<a href="#">Deployment</a>	apps/v1
<a href="#">Endpoints</a>	v1
<a href="#">EndpointSlice</a>	discovery.k8s.io/v1
<a href="#">Event</a>	v1
<a href="#">Event</a>	events.k8s.io/v1
<a href="#">Eviction</a>	policy/v1
<a href="#">FlowSchema</a>	flowcontrol.apiserver.k8s.io/v1beta3
<a href="#">HorizontalPodAutoscaler</a>	autoscaling/v2

API	API group
<a href="#">Ingress</a>	networking.k8s.io/v1
<a href="#">IngressClass</a>	networking.k8s.io/v1
<a href="#">Job</a>	batch/v1
<a href="#">Lease</a>	coordination.k8s.io/v1
<a href="#">LimitRange</a>	v1
<a href="#">LocalSubjectAccessReview</a>	authorization.k8s.io/v1
<a href="#">LogicalVolume</a>	topolvm.io/v1
<a href="#">MutatingWebhookConfiguration</a>	admissionregistration.k8s.io/v1
<a href="#">Namespace</a>	v1
<a href="#">NetworkPolicy</a>	networking.k8s.io/v1
<a href="#">Node</a>	v1
<a href="#">PersistentVolume</a>	v1
<a href="#">PersistentVolumeClaim</a>	v1
<a href="#">Pod</a>	v1
<a href="#">PodDisruptionBudget</a>	policy/v1
<a href="#">PodTemplate</a>	v1
<a href="#">PriorityClass</a>	scheduling.k8s.io/v1
<a href="#">PriorityLevelConfiguration</a>	flowcontrol.apiserver.k8s.io/v1beta3
<a href="#">RangeAllocation</a>	security.internal.openshift.io/v1
<a href="#">ReplicaSet</a>	apps/v1
<a href="#">ReplicationController</a>	v1
<a href="#">ResourceQuota</a>	v1
<a href="#">Role</a>	rbac.authorization.k8s.io/v1

API	API group
<a href="#">RoleBinding</a>	rbac.authorization.k8s.io/v1
<a href="#">Route</a>	route.openshift.io/v1
<a href="#">RuntimeClass</a>	node.k8s.io/v1
<a href="#">Scale</a>	autoscaling/v1
<a href="#">Secret</a>	v1
<a href="#">SecurityContextConstraints</a>	security.openshift.io/v1
<a href="#">SelfSubjectAccessReview</a>	authorization.k8s.io/v1
<a href="#">SelfSubjectRulesReview</a>	authorization.k8s.io/v1
<a href="#">Service</a>	v1
<a href="#">ServiceAccount</a>	v1
<a href="#">StatefulSet</a>	apps/v1
<a href="#">StorageClass</a>	storage.k8s.io/v1
<a href="#">StorageVersionMigration</a>	migration.k8s.io/v1alpha1
<a href="#">SubjectAccessReview</a>	authorization.k8s.io/v1
<a href="#">TokenRequest</a>	authentication.k8s.io/v1
<a href="#">TokenReview</a>	authentication.k8s.io/v1
<a href="#">ValidatingWebhookConfiguration</a>	admissionregistration.k8s.io/v1
<a href="#">VolumeAttachment</a>	storage.k8s.io/v1
<a href="#">VolumeSnapshot</a>	snapshot.storage.k8s.io/v1
<a href="#">VolumeSnapshotClass</a>	snapshot.storage.k8s.io/v1
<a href="#">VolumeSnapshotContent</a>	snapshot.storage.k8s.io/v1



## CHAPTER 4. API OBJECT REFERENCE

### 4.1. COMMON OBJECT REFERENCE

#### 4.1.1. io.k8s.api.admissionregistration.v1.MutatingWebhookConfigurationList schema

##### Description

MutatingWebhookConfigurationList is a list of MutatingWebhookConfiguration.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (MutatingWebhookConfiguration)</b>	List of MutatingWebhookConfiguration.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<a href="#">ListMeta</a>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.2. io.k8s.api.admissionregistration.v1.ValidatingWebhookConfigurationList schema

##### Description

ValidatingWebhookConfigurationList is a list of ValidatingWebhookConfiguration.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ValidatingWebhookConfiguration)</b>	List of ValidatingWebhookConfiguration.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<a href="#">ListMeta</a>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

### 4.1.3. io.k8s.api.apps.v1.ControllerRevisionList schema

#### Description

ControllerRevisionList is a resource containing a list of ControllerRevision objects.

#### Type

**object**

#### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ControllerRevision)</b>	Items is the list of ControllerRevisions
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<a href="#">ListMeta</a>	More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.4. io.k8s.api.apps.v1.DaemonSetList schema

##### Description

DaemonSetList is a collection of daemon sets.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (DaemonSet)</b>	A list of daemon sets.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<a href="#">ListMeta</a>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.5. io.k8s.api.apps.v1.DeploymentList schema

##### Description

DeploymentList is a list of Deployments.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<a href="#">array (Deployment)</a>	Items is the list of Deployments.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<a href="#">ListMeta</a>	Standard list metadata.

## 4.1.6. io.k8s.api.apps.v1.ReplicaSetList schema

### Description

ReplicaSetList is a collection of ReplicaSets.

### Type

**object**

### Required

- **items**

### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ReplicaSet)</b>	List of ReplicaSets. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

## 4.1.7. io.k8s.api.apps.v1.StatefulSetList schema

### Description

StatefulSetList is a collection of StatefulSets.

#### Type

**object**

#### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (StatefulSet)</b>	Items is the list of stateful sets.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

### 4.1.8. io.k8s.api.autoscaling.v2.HorizontalPodAutoscalerList schema

#### Description

HorizontalPodAutoscalerList is a list of horizontal pod autoscaler objects.

#### Type

**object**

#### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array</b> ( <a href="#">HorizontalPodAutoscaler</a> )	items is the list of horizontal pod autoscaler objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<a href="#">ListMeta</a>	metadata is the standard list metadata.

#### 4.1.9. io.k8s.api.batch.v1.CronJobList schema

##### Description

CronJobList is a collection of cron jobs.

##### Type

**object**

##### Required

- **items**

#### Schema



Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (CronJob)</b>	items is the list of CronJobs.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.10. io.k8s.api.batch.v1.JobList schema

##### Description

JobList is a collection of jobs.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Job)</b>	items is the list of Jobs.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.11. io.k8s.api.certificates.v1.CertificateSigningRequestList schema

##### Description

CertificateSigningRequestList is a collection of CertificateSigningRequest objects

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array</b> <b>(CertificateSigningRequest)</b>	items is a collection of CertificateSigningRequest objects
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	

#### 4.1.12. io.k8s.api.coordination.v1.LeaseList schema

##### Description

LeaseList is a list of Lease objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Lease)</b>	items is a list of schema objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.13. io.k8s.api.core.v1.Affinity schema

##### Description

Affinity is a group of affinity scheduling rules.

##### Type

**object**

##### Schema

Property	Type	Description
<b>nodeAffinity</b>	<b>NodeAffinity</b>	Describes node affinity scheduling rules for the pod.
<b>podAffinity</b>	<b>PodAffinity</b>	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).

Property	Type	Description
<b>podAntiAffinity</b>	<b>PodAntiAffinity</b>	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

#### 4.1.14. io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource schema

##### Description

Represents a Persistent Disk resource in AWS.

An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly value true will force the readOnly setting in VolumeMounts. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>volumeID</b>	<b>string</b>	volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>

#### 4.1.15. io.k8s.api.core.v1.AzureDiskVolumeSource schema

##### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **diskName**
- **diskURI**

##### Schema

Property	Type	Description
<b>cachingMode</b>	<b>string</b>	cachingMode is the Host Caching mode: None, Read Only, Read Write.  Possible enum values: - <b>"None"</b> - <b>"ReadOnly"</b> - <b>"ReadWrite"</b>
<b>diskName</b>	<b>string</b>	diskName is the Name of the data disk in the blob storage
<b>diskURI</b>	<b>string</b>	diskURI is the URI of data disk in the blob storage

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>kind</b>	<b>string</b>	kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared  Possible enum values: - <b>"Dedicated"</b> - <b>"Managed"</b> - <b>"Shared"</b>
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

#### 4.1.16. io.k8s.api.core.v1.AzureFilePersistentVolumeSource schema

##### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **secretName**
- **shareName**

##### Schema

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

Property	Type	Description
<b>secretName</b>	<b>string</b>	secretName is the name of secret that contains Azure Storage Account Name and Key
<b>secretNamespace</b>	<b>string</b>	secretNamespace is the namespace of the secret that contains Azure Storage Account Name and Key default is the same as the Pod
<b>shareName</b>	<b>string</b>	shareName is the azure Share Name

#### 4.1.17. io.k8s.api.core.v1.AzureFileVolumeSource schema

##### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **secretName**
- **shareName**

##### Schema

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretName</b>	<b>string</b>	secretName is the name of secret that contains Azure Storage Account Name and Key
<b>shareName</b>	<b>string</b>	shareName is the azure share Name

#### 4.1.18. io.k8s.api.core.v1.Capabilities schema

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type



**object**

## Schema

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities
<b>drop</b>	<b>array (string)</b>	Removed capabilities

**4.1.19. io.k8s.api.core.v1.CephFSPersistentVolumeSource schema**

## Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

## Type

**object**

## Required

- **monitors**

## Schema

Property	Type	Description
<b>monitors</b>	<b>array (string)</b>	monitors is Required: Monitors is a collection of Ceph monitors More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>path</b>	<b>string</b>	path is Optional: Used as the mounted root, rather than the full Ceph tree, default is /
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretFile</b>	<b>string</b>	secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

Property	Type	Description
<b>secretRef</b>	<b>SecretReference</b>	secretRef is Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>user</b>	<b>string</b>	user is Optional: User is the rados user name, default is admin More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

#### 4.1.20. io.k8s.api.core.v1.CephFSVolumeSource schema

##### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**

##### Schema

Property	Type	Description
<b>monitors</b>	<b>array (string)</b>	monitors is Required: Monitors is a collection of Ceph monitors More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>path</b>	<b>string</b>	path is Optional: Used as the mounted root, rather than the full Ceph tree, default is /
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

Property	Type	Description
<b>secretFile</b>	<b>string</b>	secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>LocalObjectReference</b>	secretRef is Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>user</b>	<b>string</b>	user is optional: User is the rados user name, default is admin More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

#### 4.1.21. io.k8s.api.core.v1.CinderPersistentVolumeSource schema

##### Description

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>secretRef</b>	<b>SecretReference</b>	secretRef is Optional: points to a secret object containing parameters used to connect to OpenStack.
<b>volumeID</b>	<b>string</b>	volumeID used to identify the volume in cinder. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

#### 4.1.22. io.k8s.api.core.v1.CinderVolumeSource schema

##### Description

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>secretRef</b>	<b>LocalObjectReference</b>	secretRef is optional: points to a secret object containing parameters used to connect to OpenStack.
<b>volumeID</b>	<b>string</b>	volumeID used to identify the volume in cinder. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

#### 4.1.23. io.k8s.api.core.v1.ClaimSource schema

##### Description

ClaimSource describes a reference to a ResourceClaim.

Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.

##### Type

**object**

##### Schema

Property	Type	Description
<b>resourceClaimName</b>	<b>string</b>	ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.

Property	Type	Description
<b>resourceClaimTemplateName</b>	<b>string</b>	<p>ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.</p> <p>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The name of the ResourceClaim will be &lt;pod name&gt;-&lt;resource name&gt;, where &lt;resource name&gt; is the PodResourceClaim.Name. Pod validation will reject the pod if the concatenated name is not valid for a ResourceClaim (e.g. too long).</p> <p>An existing ResourceClaim with that name that is not owned by the pod will not be used for the pod to avoid using an unrelated resource by mistake. Scheduling and pod startup are then blocked until the unrelated ResourceClaim is removed.</p> <p>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.</p>

#### 4.1.24. io.k8s.api.core.v1.ComponentStatusList schema

##### Description

Status of all the conditions for the component as a list of ComponentStatus objects. Deprecated: This API is deprecated in v1.19+

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ComponentStatus)</b>	List of ComponentStatus objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.25. io.k8s.api.core.v1.ConfigMapEnvSource schema

##### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

##### Type

**object**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

Property	Type	Description
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 4.1.26. io.k8s.api.core.v1.ConfigMapKeySelector schema

##### Description

Selects a key from a ConfigMap.

##### Type

**object**

##### Required

- **key**

##### Schema

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 4.1.27. io.k8s.api.core.v1.ConfigMapList schema

##### Description

ConfigMapList is a resource containing a list of ConfigMap objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ConfigMap)</b>	Items is the list of ConfigMaps.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.28. io.k8s.api.core.v1.ConfigMapProjection schema

##### Description

Adapts a ConfigMap into a projected volume.

The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>items</b>	<b>array (KeyToPath)</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

#### 4.1.29. io.k8s.api.core.v1.ConfigMapVolumeSource schema

##### Description

Adapts a ConfigMap into a volume.

The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array (KeyToPath)</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

### 4.1.30. io.k8s.api.core.v1.Container schema

#### Description

A single application container that you want to run within a pod.

#### Type

**object**

#### Required

- **name**

## Schema

Property	Type	Description
<b>args</b>	<b>array (string)</b>	Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>
<b>command</b>	<b>array (string)</b>	Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>
<b>env</b>	<b>array (EnvVar)</b>	List of environment variables to set in the container. Cannot be updated.

Property	Type	Description
<b>envFrom</b>	<b>array (EnvFromSource)</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
<b>imagePullPolicy</b>	<b>string</b>	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a>  Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present
<b>lifecycle</b>	<b>Lifecycle</b>	Actions that the management system should take in response to container lifecycle events. Cannot be updated.

Property	Type	Description
<b>livenessProbe</b>	<b>Probe</b>	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>name</b>	<b>string</b>	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
<b>ports</b>	<b>array (ContainerPort)</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>readinessProbe</b>	<b>Probe</b>	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>resizePolicy</b>	<b>array (ContainerResizePolicy)</b>	Resources resize policy for the container.
<b>resources</b>	<b>ResourceRequirements</b>	Compute Resources required by this container. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>securityContext</b>	<b>SecurityContext</b>	SecurityContext defines the security options the container should be run with. If set, the fields of SecurityContext override the equivalent fields of PodSecurityContext. More info: <a href="https://kubernetes.io/docs/tasks/configure-pod-container/security-context/">https://kubernetes.io/docs/tasks/configure-pod-container/security-context/</a>
<b>startupProbe</b>	<b>Probe</b>	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

Property	Type	Description
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.



Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array (VolumeDevice)</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeMounts</b>	<b>array (VolumeMount)</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

## 4.1.31. io.k8s.api.core.v1.ContainerPort schema

## Description

ContainerPort represents a network port in a single container.

## Type

**object**

## Required

- **containerPort**

## Schema

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.

### 4.1.32. io.k8s.api.core.v1.ContainerResizePolicy schema

## Description

ContainerResizePolicy represents resource resize policy for the container.

## Type

**object**

## Required

- **resourceName**

- **restartPolicy**

#### Schema

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

### 4.1.33. io.k8s.api.core.v1.CSIPersistentVolumeSource schema

#### Description

Represents storage that is managed by an external CSI volume driver (Beta feature)

#### Type

**object**

#### Required

- **driver**
- **volumeHandle**

#### Schema

Property	Type	Description
<b>controllerExpandSecretRef</b>	<b>SecretReference</b>	controllerExpandSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI ControllerExpandVolume call. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secrets are passed.

Property	Type	Description
<b>controllerPublishSecretRef</b>	<b>SecretReference</b>	controllerPublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI ControllerPublishVolume and ControllerUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secrets are passed.
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume. Required.
<b>fsType</b>	<b>string</b>	fsType to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs".
<b>nodeExpandSecretRef</b>	<b>SecretReference</b>	nodeExpandSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodeExpandVolume call. This is a beta field which is enabled default by CSINodeExpandSecret feature gate. This field is optional, may be omitted if no secret is required. If the secret object contains more than one secret, all secrets are passed.
<b>nodePublishSecretRef</b>	<b>SecretReference</b>	nodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secrets are passed.

Property	Type	Description
<b>nodeStageSecretRef</b>	<b>SecretReference</b>	nodeStageSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodeStageVolume and NodeStageVolume and NodeUnstageVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secrets are passed.
<b>readOnly</b>	<b>boolean</b>	readOnly value to pass to ControllerPublishVolumeRequest. Defaults to false (read/write).
<b>volumeAttributes</b>	<b>object (string)</b>	volumeAttributes of the volume to publish.
<b>volumeHandle</b>	<b>string</b>	volumeHandle is the unique volume name returned by the CSI volume plugin's CreateVolume to refer to the volume on all subsequent calls. Required.

#### 4.1.34. io.k8s.api.core.v1.CSIVolumeSource schema

##### Description

Represents a source location of a volume to mount, managed by an external CSI driver

##### Type

**object**

##### Required

- **driver**

##### Schema

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
<b>nodePublishSecretRef</b>	<b>LocalObjectReference</b>	nodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
<b>readOnly</b>	<b>boolean</b>	readOnly specifies a read-only configuration for the volume. Defaults to false (read/write).
<b>volumeAttributes</b>	<b>object (string)</b>	volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

#### 4.1.35. io.k8s.api.core.v1.DownwardAPIProjection schema

##### Description

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

##### Type

**object**

##### Schema

Property	Type	Description
<b>items</b>	<b>array (DownwardAPIVolumeFile)</b>	Items is a list of DownwardAPIVolume file

#### 4.1.36. io.k8s.api.core.v1.DownwardAPIVolumeFile schema

##### Description

DownwardAPIVolumeFile represents information to create the file containing the pod field

Type

**object**

Required

- **path**

Schema

Property	Type	Description
<b>fieldRef</b>	<b>ObjectFieldSelector</b>	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
<b>resourceFieldRef</b>	<b>ResourceFieldSelector</b>	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

#### 4.1.37. io.k8s.api.core.v1.DownwardAPIVolumeSource schema

Description

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

Type

**object**

## Schema

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b> ( <a href="#">DownwardAPIVolumeFile</a> )	Items is a list of downward API volume file

## 4.1.38. io.k8s.api.core.v1.EmptyDirVolumeSource schema

## Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

## Type

**object**

## Schema

Property	Type	Description
<b>medium</b>	<b>string</b>	medium represents what type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>



Property	Type	Description
<b>sizeLimit</b>	<b>Quantity</b>	sizeLimit is the total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

### 4.1.39. io.k8s.api.core.v1.EndpointsList schema

#### Description

EndpointsList is a list of endpoints.

#### Type

**object**

#### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Endpoints)</b>	List of endpoints.

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.40. io.k8s.api.core.v1.EnvFromSource schema

##### Description

EnvFromSource represents the source of a set of ConfigMaps

##### Type

**object**

##### Schema

Property	Type	Description
<b>configMapRef</b>	<b>ConfigMapEnvSource</b>	The ConfigMap to select from
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>SecretEnvSource</b>	The Secret to select from

#### 4.1.41. io.k8s.api.core.v1.EnvVar schema

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

## Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>EnvVarSource</b>	Source for the environment variable's value. Cannot be used if value is not empty.

## 4.1.42. io.k8s.api.core.v1.EnvVarSource schema

## Description

EnvVarSource represents a source for the value of an EnvVar.

## Type

**object**

## Schema

Property	Type	Description
<b>configMapKeyRef</b>	<b>ConfigMapKeySelector</b>	Selects a key of a ConfigMap.

Property	Type	Description
<b>fieldRef</b>	<b>ObjectFieldSelector</b>	Selects a field of the pod: supports metadata.name, metadata.namespace, <b>metadata.labels['&lt;KEY&gt;']</b> , <b>metadata.annotations['&lt;KEY&gt;']</b> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
<b>resourceFieldRef</b>	<b>ResourceFieldSelector</b>	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
<b>secretKeyRef</b>	<b>SecretKeySelector</b>	Selects a key of a secret in the pod's namespace

#### 4.1.43. io.k8s.api.core.v1.EphemeralContainer schema

##### Description

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

##### Type

**object**

##### Required

- **name**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Endpoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array (EnvVar)</b>	<p>List of environment variables to set in the container. Cannot be updated.</p>

Property	Type	Description
<b>envFrom</b>	<b>array (EnvFromSource)</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a>
<b>imagePullPolicy</b>	<b>string</b>	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a>  Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail If the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present
<b>lifecycle</b>	<b>Lifecycle</b>	Lifecycle is not allowed for ephemeral containers.
<b>livenessProbe</b>	<b>Probe</b>	Probes are not allowed for ephemeral containers.

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
<b>ports</b>	<b>array (ContainerPort)</b>	Ports are not allowed for ephemeral containers.
<b>readinessProbe</b>	<b>Probe</b>	Probes are not allowed for ephemeral containers.
<b>resizePolicy</b>	<b>array (ContainerResizePolicy)</b>	Resources resize policy for the container.
<b>resources</b>	<b>ResourceRequirements</b>	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
<b>securityContext</b>	<b>SecurityContext</b>	Optional: SecurityContext defines the security options the ephemeral container should be run with. If set, the fields of SecurityContext override the equivalent fields of PodSecurityContext.
<b>startupProbe</b>	<b>Probe</b>	Probes are not allowed for ephemeral containers.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

Property	Type	Description
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>targetContainerName</b>	<b>string</b>	<p>If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.</p> <p>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined.</p>
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.



Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array (VolumeDevice)</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeMounts</b>	<b>array (VolumeMount)</b>	Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

Property	Type	Description
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 4.1.44. io.k8s.api.core.v1.EphemeralVolumeSource schema

##### Description

Represents an ephemeral volume that is handled by a normal storage driver.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>volumeClaimTemplate</b>	<b>PersistentVolumeClaimTemplate</b>	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <b>&lt;pod name&gt;-&lt;volume name&gt;</b> where <b>&lt;volume name&gt;</b> is the name from the <b>PodSpec.Volumes</b> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <b>not</b> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

#### 4.1.45. io.k8s.api.core.v1.EventList schema

##### Description

EventList is a list of events.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Event)</b>	List of events
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.46. io.k8s.api.core.v1.EventSource schema

##### Description

EventSource contains information for an event.

##### Type

**object**

##### Schema

Property	Type	Description
<b>component</b>	<b>string</b>	Component from which the event is generated.
<b>host</b>	<b>string</b>	Node name on which the event is generated.

#### 4.1.47. io.k8s.api.core.v1.ExecAction schema

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

##### Schema

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 4.1.48. io.k8s.api.core.v1.FCVolumeSource schema

##### Description

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>lun</b>	<b>integer</b>	lun is Optional: FC target lun number
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

Property	Type	Description
<b>targetWWNs</b>	<b>array (string)</b>	targetWWNs is Optional: FC target worldwide names (WWNs)
<b>wwids</b>	<b>array (string)</b>	wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

#### 4.1.49. io.k8s.api.core.v1.FlexPersistentVolumeSource schema

##### Description

FlexPersistentVolumeSource represents a generic persistent volume resource that is provisioned/attached using an exec based plugin.

##### Type

**object**

##### Required

- **driver**

##### Schema

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume.
<b>fsType</b>	<b>string</b>	fsType is the Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
<b>options</b>	<b>object (string)</b>	options is Optional: this field holds extra command options if any.
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

Property	Type	Description
<b>secretRef</b>	<a href="#">SecretReference</a>	secretRef is Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

#### 4.1.50. io.k8s.api.core.v1.FlexVolumeSource schema

##### Description

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

##### Type

**object**

##### Required

- **driver**

##### Schema

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume.
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
<b>options</b>	<b>object (string)</b>	options is Optional: this field holds extra command options if any.
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

Property	Type	Description
<b>secretRef</b>	<a href="#">LocalObjectReference</a>	secretRef is Optional: secretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

#### 4.1.51. io.k8s.api.core.v1.FlockerVolumeSource schema

##### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Schema

Property	Type	Description
<b>datasetName</b>	<b>string</b>	datasetName is Name of the dataset stored as metadata → name on the dataset for Flocker should be considered as deprecated
<b>datasetUUID</b>	<b>string</b>	datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset

#### 4.1.52. io.k8s.api.core.v1.GCEPersistentDiskVolumeSource schema

##### Description

Represents a Persistent Disk resource in Google Compute Engine.

A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **pdName**



## Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>pdName</b>	<b>string</b>	pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

## 4.1.53. io.k8s.api.core.v1.GitRepoVolumeSource schema

## Description

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.

DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

**Type****object****Required**

- **repository**

**Schema**

Property	Type	Description
<b>directory</b>	<b>string</b>	directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
<b>repository</b>	<b>string</b>	repository is the URL
<b>revision</b>	<b>string</b>	revision is the commit hash for the specified revision.

**4.1.54. io.k8s.api.core.v1.GlusterfsPersistentVolumeSource schema****Description**

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

**Type****object****Required**

- **endpoints**
- **path**

**Schema**

Property	Type	Description
<b>endpoints</b>	<b>string</b>	endpoints is the endpoint name that details Glusterfs topology. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

Property	Type	Description
<b>endpointsNamespace</b>	<b>string</b>	endpointsNamespace is the namespace that contains Glusterfs endpoint. If this field is empty, the EndpointNamespace defaults to the same namespace as the bound PVC. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>path</b>	<b>string</b>	path is the Glusterfs volume path. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

#### 4.1.55. io.k8s.api.core.v1.GlusterfsVolumeSource schema

##### Description

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **endpoints**
- **path**

##### Schema

Property	Type	Description
<b>endpoints</b>	<b>string</b>	endpoints is the endpoint name that details Glusterfs topology. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

Property	Type	Description
<b>path</b>	<b>string</b>	path is the Glusterfs volume path. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

#### 4.1.56. io.k8s.api.core.v1.GRPCAction schema

##### Description

##### Type

**object**

##### Required

- **port**

##### Schema

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 4.1.57. io.k8s.api.core.v1.HostAlias schema

##### Description

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

Type

**object**

Schema

Property	Type	Description
<b>hostnames</b>	<b>array (string)</b>	Hostnames for the above IP address.
<b>ip</b>	<b>string</b>	IP address of the host file entry.

#### 4.1.58. io.k8s.api.core.v1.HostPathVolumeSource schema

Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

Type

**object**

Required

- **path**

Schema

Property	Type	Description
<b>path</b>	<b>string</b>	path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>

Property	Type	Description
<b>type</b>	<b>string</b>	<p>type for HostPath Volume Defaults to "" More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a></p> <p>Possible enum values: - "" For backwards compatible, leave it empty if unset - <b>"BlockDevice"</b> A block device must exist at the given path - <b>"CharDevice"</b> A character device must exist at the given path - <b>"Directory"</b> A directory must exist at the given path - <b>"DirectoryOrCreate"</b> If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet. - <b>"File"</b> A file must exist at the given path - <b>"FileOrCreate"</b> If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet. - <b>"Socket"</b> A UNIX socket must exist at the given path</p>

#### 4.1.59. io.k8s.api.core.v1.HTTPGetAction schema

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

##### Schema

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array (HTTPHeader)</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 4.1.60. io.k8s.api.core.v1.HTTPHeader schema

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 4.1.61. io.k8s.api.core.v1.ISCSIPersistentVolumeSource schema

## Description

ISCSIPersistentVolumeSource represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.

## Type

**object**

## Required

- **targetPortal**
- **iqn**
- **lun**

## Schema

Property	Type	Description
<b>chapAuthDiscovery</b>	<b>boolean</b>	chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication
<b>chapAuthSession</b>	<b>boolean</b>	chapAuthSession defines whether support iSCSI Session CHAP authentication
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#iscsi">https://kubernetes.io/docs/concepts/storage/volumes#iscsi</a>
<b>initiatorName</b>	<b>string</b>	initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <target portal>:<volume name> will be created for the connection.
<b>iqn</b>	<b>string</b>	iqn is Target iSCSI Qualified Name.
<b>iscsiInterface</b>	<b>string</b>	iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).



Property	Type	Description
<b>lun</b>	<b>integer</b>	lun is iSCSI Target Lun number.
<b>portals</b>	<b>array (string)</b>	portals is the iSCSI Target Portal List. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
<b>secretRef</b>	<b>SecretReference</b>	secretRef is the CHAP Secret for iSCSI target and initiator authentication
<b>targetPortal</b>	<b>string</b>	targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

#### 4.1.62. io.k8s.api.core.v1.ISCSIVolumeSource schema

##### Description

Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **targetPortal**
- **iqn**
- **lun**

##### Schema

Property	Type	Description
<b>chapAuthDiscovery</b>	<b>boolean</b>	chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication

Property	Type	Description
<b>chapAuthSession</b>	<b>boolean</b>	chapAuthSession defines whether support iSCSI Session CHAP authentication
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#iscsi">https://kubernetes.io/docs/concepts/storage/volumes#iscsi</a>
<b>initiatorName</b>	<b>string</b>	initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <target portal>:<volume name> will be created for the connection.
<b>iqn</b>	<b>string</b>	iqn is the target iSCSI Qualified Name.
<b>iscsilInterface</b>	<b>string</b>	iscsilInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
<b>lun</b>	<b>integer</b>	lun represents iSCSI Target Lun number.
<b>portals</b>	<b>array (string)</b>	portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
<b>secretRef</b>	<b>LocalObjectReference</b>	secretRef is the CHAP Secret for iSCSI target and initiator authentication

Property	Type	Description
<b>targetPortal</b>	<b>string</b>	targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

#### 4.1.63. io.k8s.api.core.v1.KeyToPath schema

##### Description

Maps a string key to a path within a volume.

##### Type

**object**

##### Required

- **key**
- **path**

##### Schema

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 4.1.64. io.k8s.api.core.v1.Lifecycle schema

##### Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

## Type

**object**

## Schema

Property	Type	Description
<b>postStart</b>	<a href="#">LifecycleHandler</a>	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: <a href="https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks">https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks</a>
<b>preStop</b>	<a href="#">LifecycleHandler</a>	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The Pod's termination grace period countdown begins before the PreStop hook is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period (unless delayed by finalizers). Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: <a href="https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks">https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks</a>

### 4.1.65. io.k8s.api.core.v1.LifecycleHandler schema

#### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## Type

**object**

## Schema

Property	Type	Description
<b>exec</b>	<a href="#">ExecAction</a>	Exec specifies the action to take.
<b>httpGet</b>	<a href="#">HTTPGetAction</a>	HTTPGet specifies the http request to perform.
<b>tcpSocket</b>	<a href="#">TCPSocketAction</a>	Deprecated. TCPSocket is NOT supported as a LifecycleHandler and kept for the backward compatibility. There are no validation of this field and lifecycle hooks will fail in runtime when tcp handler is specified.

### 4.1.66. io.k8s.api.core.v1.LimitRangeList schema

## Description

LimitRangeList is a list of LimitRange items.

## Type

**object**

## Required

- **items**

## Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>items</b>	<b>array (LimitRange)</b>	Items is a list of LimitRange objects. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.67. io.k8s.api.core.v1.LocalObjectReference schema

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 4.1.68. io.k8s.api.core.v1.LocalVolumeSource schema

##### Description

Local represents directly-attached storage with node affinity (Beta feature)

##### Type

**object**

##### Required

- **path**

#### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. It applies only when the Path is a block device. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default value is to auto-select a filesystem if unspecified.
<b>path</b>	<b>string</b>	path of the full path to the volume on the node. It can be either a directory or block device (disk, partition, ...).

#### 4.1.69. io.k8s.api.core.v1.NamespaceList schema

##### Description

NamespaceList is a list of Namespaces.

##### Type

**object**

##### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Namespace)</b>	Items is the list of Namespace objects in the list. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.70. io.k8s.api.core.v1.NFSVolumeSource schema

##### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **server**
- **path**

##### Schema

Property	Type	Description
<b>path</b>	<b>string</b>	path that is exported by the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>



Property	Type	Description
<b>server</b>	<b>string</b>	server is the hostname or IP address of the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

#### 4.1.71. io.k8s.api.core.v1.NodeAffinity schema

##### Description

Node affinity is a group of node affinity scheduling rules.

##### Type

**object**

##### Schema

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array (PreferredSchedulingTerm)</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>NodeSelector</b>	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

## 4.1.72. io.k8s.api.core.v1.NodeList schema

### Description

NodeList is the whole list of all Nodes which have been registered with master.

### Type

**object**

### Required

- **items**

### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Node)</b>	List of nodes
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

## 4.1.73. io.k8s.api.core.v1.NodeSelector schema

### Description

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

### Type

**object**

## Required

- **nodeSelectorTerms**

## Schema

Property	Type	Description
<b>nodeSelectorTerms</b>	<b>array (NodeSelectorTerm)</b>	Required. A list of node selector terms. The terms are ORed.

## 4.1.74. io.k8s.api.core.v1.NodeSelectorRequirement schema

## Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

## Type

**object**

## Required

- **key**
- **operator**

## Schema

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>

Property	Type	Description
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 4.1.75. io.k8s.api.core.v1.NodeSelectorTerm schema

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

##### Schema

Property	Type	Description
<b>matchExpressions</b>	<b>array (NodeSelectorRequirement)</b>	A list of node selector requirements by node's labels.
<b>matchFields</b>	<b>array (NodeSelectorRequirement)</b>	A list of node selector requirements by node's fields.

#### 4.1.76. io.k8s.api.core.v1.ObjectFieldSelector schema

##### Description

ObjectFieldSelector selects an APIVersioned field of an object.

##### Type

**object**

##### Required

- **fieldPath**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 4.1.77. io.k8s.api.core.v1.ObjectReference schema

##### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

##### Type

**object**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

#### 4.1.78. io.k8s.api.core.v1.PersistentVolumeClaim schema

##### Description

PersistentVolumeClaim is a user's request for and claim to a persistent volume

##### Type

**object**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes
<b>status</b>	<b>object</b>	PersistentVolumeClaimStatus is the current status of a persistent volume claim.

..spec Description::

+

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the desired access modes the volume should have. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>

Property	Type	Description
<b>dataSource</b>	<b>object</b>	TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.



Property	Type	Description
<b>dataSourceRef</b>	<b>object</b>	<p>dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: * While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. * While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. * While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.</p>

Property	Type	Description
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over volumes to consider for binding.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of the StorageClass required by the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1</a>
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>
<b>volumeName</b>	<b>string</b>	volumeName is the binding reference to the PersistentVolume backing this claim.

..spec.dataSource Description::

+

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

### Type

**object**

Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

..spec.dataSourceRef Description::

+

dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: \* While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. \* While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. \* While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

## Type

**object**

### Required

- **kind**
- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced
<b>namespace</b>	<b>string</b>	Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

..spec.resources Description::

+

ResourceRequirements describes the compute resource requirements.

**Type**

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in <code>spec.resourceClaims</code>, that are used by this container.</p> <p>This is an alpha field and requires enabling the <code>DynamicResourceAllocation</code> feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in <code>PodSpec.ResourceClaims</code> .
<b>limits</b>	<b>object (Quantity)</b>	<p>Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a></p>
<b>requests</b>	<b>object (Quantity)</b>	<p>Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a></p>

`..spec.resources.claims` Description::

+

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container.

This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

#### Type

**array**

`..spec.resources.claims[]` Description::

+

ResourceClaim references one entry in PodSpec.ResourceClaims.

### Type

**object**

### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container.

..status Description::

+

PersistentVolumeClaimStatus is the current status of a persistent volume claim.

### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the actual access modes the volume backing the PVC has. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>

Property	Type	Description
<b>allocatedResources</b>	<b>object (Quantity)</b>	allocatedResources is the storage resource within AllocatedResources tracks the capacity allocated to a PVC. It may be larger than the actual capacity when a volume expansion operation is requested. For storage quota, the larger value from allocatedResources and PVC.spec.resources is used. If allocatedResources is not set, PVC.spec.resources alone is used for quota calculation. If a volume expansion capacity request is lowered, allocatedResources is only lowered if there are no expansion operations in progress and if the actual volume capacity is equal or lower than the requested capacity. This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature.
<b>capacity</b>	<b>object (Quantity)</b>	capacity represents the actual resources of the underlying volume.
<b>conditions</b>	<b>array</b>	conditions is the current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
<b>conditions[]</b>	<b>object</b>	PersistentVolumeClaimCondition contains details about state of pvc

Property	Type	Description
<b>phase</b>	<b>string</b>	<p>phase represents the current phase of PersistentVolumeClaim.</p> <p>Possible enum values: - <b>"Bound"</b> used for PersistentVolumeClaims that are bound - <b>"Lost"</b> used for PersistentVolumeClaims that lost their underlying PersistentVolume. The claim was bound to a PersistentVolume and this volume does not exist any longer and all data on it was lost. - <b>"Pending"</b> used for PersistentVolumeClaims that are not yet bound</p>



Property	Type	Description
<b>resizeStatus</b>	<b>string</b>	<p>resizeStatus stores status of resize operation. ResizeStatus is not set by default but when expansion is complete resizeStatus is set to empty string by resize controller or kubelet. This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature.</p> <p>Possible enum values: - "" When expansion is complete, the empty string is set by resize controller or kubelet. - <b>"ControllerExpansionFailed"</b> State set when expansion has failed in resize controller with a terminal error. Transient errors such as timeout should not set this status and should leave ResizeStatus unmodified, so as resize controller can resume the volume expansion. - <b>"ControllerExpansionInProgress"</b> State set when resize controller starts expanding the volume in control-plane - <b>"NodeExpansionFailed"</b> State set when expansion has failed in kubelet with a terminal error. Transient errors don't set NodeExpansionFailed. - <b>"NodeExpansionInProgress"</b> State set when kubelet starts expanding the volume. - <b>"NodeExpansionPending"</b> State set when resize controller has finished expanding the volume but further expansion is needed on the node.</p>

..status.conditions Description::

+

conditions is the current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.

Type

**array**

..status.conditions[] Description::

+

PersistentVolumeClaimCondition contains details about state of pvc

Type

**object**

Required

- **type**
- **status**

Property	Type	Description
<b>lastProbeTime</b>	<b>Time</b>	lastProbeTime is the time we probed the condition.
<b>lastTransitionTime</b>	<b>Time</b>	lastTransitionTime is the time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	message is the human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	reason is a unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
<b>status</b>	<b>string</b>	
<b>type</b>	<b>string</b>	

#### 4.1.79. io.k8s.api.core.v1.PersistentVolumeClaimList schema

Description

PersistentVolumeClaimList is a list of PersistentVolumeClaim items.

Type

**object**

Required

- **items**

## Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array</b> <b>(PersistentVolumeClaim)</b>	items is a list of persistent volume claims. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims">https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

## 4.1.80. io.k8s.api.core.v1.PersistentVolumeClaimSpec schema

## Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

## Type

**object**

## Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the desired access modes the volume should have. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>
<b>dataSource</b>	<b>TypedLocalObjectReference</b>	dataSource field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot) * An existing PVC (PersistentVolumeClaim) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. When the AnyVolumeDataSource feature gate is enabled, dataSource contents will be copied to dataSourceRef, and dataSourceRef contents will be copied to dataSource when dataSourceRef.namespace is not specified. If the namespace is specified, then dataSourceRef will not be copied to dataSource.

Property	Type	Description
<b>dataSourceRef</b>	<b>TypedObjectReference</b>	<p>dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: * While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. * While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. * While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.</p>

Property	Type	Description
<b>resources</b>	<b>ResourceRequirements</b>	resources represents the minimum resources the volume should have. If RecoverVolumeExpansionFailure feature is enabled users are allowed to specify resource requirements that are lower than previous value but must still be higher than capacity recorded in the status field of the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources">https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources</a>
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over volumes to consider for binding.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of the StorageClass required by the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1</a>
<b>volumeMode</b>	<b>string</b>	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.  Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.
<b>volumeName</b>	<b>string</b>	volumeName is the binding reference to the PersistentVolume backing this claim.

#### 4.1.81. io.k8s.api.core.v1.PersistentVolumeClaimTemplate schema

##### Description

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

##### Type

**object**

**Required**

- **spec**

**Schema**

Property	Type	Description
<b>metadata</b>	<a href="#">ObjectMeta</a>	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
<b>spec</b>	<a href="#">PersistentVolumeClaimSpec</a>	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

**4.1.82. io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource schema****Description**

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A

PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

**Type**

**object**

**Required**

- **claimName**

**Schema**

Property	Type	Description
<b>claimName</b>	<b>string</b>	claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims">https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims</a>
<b>readOnly</b>	<b>boolean</b>	readOnly Will force the ReadOnly setting in VolumeMounts. Default false.

### 4.1.83. io.k8s.api.core.v1.PersistentVolumeList schema

#### Description

PersistentVolumeList is a list of PersistentVolume items.

#### Type

**object**

#### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (PersistentVolume)</b>	items is a list of persistent volumes. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes">https://kubernetes.io/docs/concepts/storage/persistent-volumes</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

### 4.1.84. io.k8s.api.core.v1.PersistentVolumeSpec schema

#### Description

PersistentVolumeSpec is the specification of a persistent volume.



Type  
object

Schema

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains all ways the volume can be mounted. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes</a>
<b>awsElasticBlockStore</b>	<b><a href="#">AWSElasticBlockStoreVolumeSource</a></b>	awsElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>azureDisk</b>	<b><a href="#">AzureDiskVolumeSource</a></b>	azureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
<b>azureFile</b>	<b><a href="#">AzureFilePersistentVolumeSource</a></b>	azureFile represents an Azure File Service mount on the host and bind mount to the pod.
<b>capacity</b>	<b>object (Quantity)</b>	capacity is the description of the persistent volume's resources and capacity. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#capacity">https://kubernetes.io/docs/concepts/storage/persistent-volumes#capacity</a>
<b>cephfs</b>	<b><a href="#">CephFSPersistentVolumeSource</a></b>	cephFS represents a Ceph FS mount on the host that shares a pod's lifetime
<b>cinder</b>	<b><a href="#">CinderPersistentVolumeSource</a></b>	cinder represents a cinder volume attached and mounted on kubelets host machine. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

Property	Type	Description
<b>claimRef</b>	<b>ObjectReference</b>	claimRef is part of a bi-directional binding between PersistentVolume and PersistentVolumeClaim. Expected to be non-nil when bound. claim.VolumeName is the authoritative bind between PV and PVC. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#binding">https://kubernetes.io/docs/concepts/storage/persistent-volumes#binding</a>
<b>csi</b>	<b>CSIPersistentVolumeSource</b>	csi represents storage that is handled by an external CSI driver (Beta feature).
<b>fc</b>	<b>FCVolumeSource</b>	fc represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
<b>flexVolume</b>	<b>FlexPersistentVolumeSource</b>	flexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
<b>flocker</b>	<b>FlockerVolumeSource</b>	flocker represents a Flocker volume attached to a kubelet's host machine and exposed to the pod for its usage. This depends on the Flocker control service being running
<b>gcePersistentDisk</b>	<b>GCEPersistentDiskVolumeSource</b>	gcePersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. Provisioned by an admin. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>glusterfs</b>	<b>GlusterfsPersistentVolumeSource</b>	glusterfs represents a Glusterfs volume that is attached to a host and exposed to the pod. Provisioned by an admin. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md">https://examples.k8s.io/volumes/glusterfs/README.md</a>

Property	Type	Description
<b>hostPath</b>	<b>HostPathVolumeSource</b>	hostPath represents a directory on the host. Provisioned by a developer or tester. This is useful for single-node development and testing only! On-host storage is not supported in any way and WILL NOT WORK in a multi-node cluster. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>
<b>iscsi</b>	<b>ISCSIPersistentVolumeSource</b>	iscsi represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. Provisioned by an admin.
<b>local</b>	<b>LocalVolumeSource</b>	local represents directly-attached storage with node affinity
<b>mountOptions</b>	<b>array (string)</b>	mountOptions is the list of mount options, e.g. ["ro", "soft"]. Not validated - mount will simply fail if one is invalid. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes/#mount-options">https://kubernetes.io/docs/concepts/storage/persistent-volumes/#mount-options</a>
<b>nfs</b>	<b>NFSVolumeSource</b>	nfs represents an NFS mount on the host. Provisioned by an admin. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>nodeAffinity</b>	<b>VolumeNodeAffinity</b>	nodeAffinity defines constraints that limit what nodes this volume can be accessed from. This field influences the scheduling of pods that use this volume.

Property	Type	Description
<b>persistentVolumeReclaimPolicy</b>	<b>string</b>	<p>persistentVolumeReclaimPolicy defines what happens to a persistent volume when released from its claim. Valid options are Retain (default for manually created PersistentVolumes), Delete (default for dynamically provisioned PersistentVolumes), and Recycle (deprecated). Recycle must be supported by the volume plugin underlying this PersistentVolume. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#reclaiming">https://kubernetes.io/docs/concepts/storage/persistent-volumes#reclaiming</a></p> <p>Possible enum values: - <b>"Delete"</b> means the volume will be deleted from Kubernetes on release from its claim. The volume plugin must support Deletion. - <b>"Recycle"</b> means the volume will be recycled back into the pool of unbound persistent volumes on release from its claim. The volume plugin must support Recycling. - <b>"Retain"</b> means the volume will be left in its current phase (Released) for manual reclamation by the administrator. The default policy is Retain.</p>
<b>photonPersistentDisk</b>	<b>PhotonPersistentDiskVolumeSource</b>	photonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
<b>portworxVolume</b>	<b>PortworxVolumeSource</b>	portworxVolume represents a portworx volume attached and mounted on kubelets host machine
<b>quobyte</b>	<b>QuobyteVolumeSource</b>	quobyte represents a Quobyte mount on the host that shares a pod's lifetime
<b>rbd</b>	<b>RBDPersistentVolumeSource</b>	rbd represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md">https://examples.k8s.io/volumes/rbd/README.md</a>

Property	Type	Description
<b>scaleIO</b>	<a href="#">ScaleIOPersistentVolumeSource</a>	scaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of StorageClass to which this persistent volume belongs. Empty value means that this volume does not belong to any StorageClass.
<b>storageos</b>	<a href="#">StorageOSPersistentVolumeSource</a>	storageOS represents a StorageOS volume that is attached to the kubelet's host machine and mounted into the pod More info: <a href="https://examples.k8s.io/volumes/storageos/README.md">https://examples.k8s.io/volumes/storageos/README.md</a>
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines if a volume is intended to be used with a formatted filesystem or to remain in raw block state. Value of Filesystem is implied when not included in spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>
<b>vsphereVolume</b>	<a href="#">VsphereVirtualDiskVolumeSource</a>	vsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

#### 4.1.85. io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource schema

##### Description

Represents a Photon Controller persistent disk resource.

##### Type

**object**

## Required

- **pdID**

## Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>pdID</b>	<b>string</b>	pdID is the ID that identifies Photon Controller persistent disk

**4.1.86. io.k8s.api.core.v1.PodAffinity schema**

## Description

Pod affinity is a group of inter pod affinity scheduling rules.

## Type

**object**

## Schema

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b> <b>(WeightedPodAffinityTerm)</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array (PodAffinityTerm)</b>	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

#### 4.1.87. io.k8s.api.core.v1.PodAffinityTerm schema

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

##### Schema

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} matches all namespaces.

Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 4.1.88. io.k8s.api.core.v1.PodAntiAffinity schema

##### Description

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array (WeightedPodAffinityTerm)</b>	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array (PodAffinityTerm)</b>	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

#### 4.1.89. io.k8s.api.core.v1.PodDNSConfig schema

##### Description

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

##### Type

**object**

##### Schema

Property	Type	Description
<b>nameservers</b>	<b>array (string)</b>	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
<b>options</b>	<b>array (PodDNSConfigOption)</b>	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
<b>searches</b>	<b>array (string)</b>	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

#### 4.1.90. io.k8s.api.core.v1.PodDNSConfigOption schema

##### Description

PodDNSConfigOption defines DNS resolver options of a pod.

##### Type

**object**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Required.
<b>value</b>	<b>string</b>	

#### 4.1.91. io.k8s.api.core.v1.PodList schema

##### Description

PodList is a list of Pods.

##### Type

**object**

##### Required

- **items**

## Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Pod)</b>	List of pods. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

## 4.1.92. io.k8s.api.core.v1.PodOS schema

## Description

PodOS defines the OS parameters of a pod.

## Type

**object**

## Required

- **name**

## Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: <a href="https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration">https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration</a> Clients should expect to handle additional values and treat unrecognized values in this field as os: null

#### 4.1.93. io.k8s.api.core.v1.PodReadinessGate schema

##### Description

PodReadinessGate contains the reference to a pod condition

##### Type

**object**

##### Required

- **conditionType**

##### Schema

Property	Type	Description
<b>conditionType</b>	<b>string</b>	ConditionType refers to a condition in the pod's condition list with matching type.

#### 4.1.94. io.k8s.api.core.v1.PodResourceClaim schema

##### Description

PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

##### Type

**object**

##### Required

- **name**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL.
<b>source</b>	<a href="#">ClaimSource</a>	Source describes where to find the ResourceClaim.

#### 4.1.95. io.k8s.api.core.v1.PodSchedulingGate schema

##### Description

PodSchedulingGate is associated to a Pod to guard its scheduling.

##### Type

**object**

##### Required

- **name**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the scheduling gate. Each scheduling gate must have a unique name field.

#### 4.1.96. io.k8s.api.core.v1.PodSecurityContext schema

##### Description

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>fsGroup</b>	<b>integer</b>	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none"> <li>1. The owning GID will be the FSGroup</li> <li>2. The setgid bit is set (new files created in the volume will be owned by FSGroup)</li> <li>3. The permission bits are OR'd with rw-rw----</li> </ol> <p>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows.</p>
<b>fsGroupChangePolicy</b>	<b>string</b>	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership(and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Always"</b> indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior. -  <b>"OnRootMismatch"</b> indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume.</p>

Property	Type	Description
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>seLinuxOptions</b>	<b>SELinuxOptions</b>	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.
<b>seccompProfile</b>	<b>SeccompProfile</b>	The seccomp options to use by the containers in this pod. Note that this field cannot be set when spec.os.name is windows.
<b>supplementalGroups</b>	<b>array (integer)</b>	A list of groups applied to the first process run in each container, in addition to the container's primary GID, the fsGroup (if specified), and group memberships defined in the container image for the uid of the container process. If unspecified, no additional groups are added to any container. Note that group memberships defined in the container image for the uid of the container process are still effective, even if they are not included in this list. Note that this field cannot be set when spec.os.name is windows.
<b>sysctls</b>	<b>array (Sysctl)</b>	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.



Property	Type	Description
<b>windowsOptions</b>	<b>WindowsSecurityContextOptions</b>	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is linux.

#### 4.1.97. io.k8s.api.core.v1.PodSpec schema

##### Description

PodSpec is a description of a pod.

##### Type

**object**

##### Required

- **containers**

##### Schema

Property	Type	Description
<b>activeDeadlineSeconds</b>	<b>integer</b>	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
<b>affinity</b>	<b>Affinity</b>	If specified, the pod's scheduling constraints
<b>automountServiceAccountToken</b>	<b>boolean</b>	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
<b>containers</b>	<b>array (Container)</b>	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

Property	Type	Description
<b>dnsConfig</b>	<b>PodDNSConfig</b>	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
<b>dnsPolicy</b>	<b>string</b>	<p>Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.</p> <p>Possible enum values: - <b>"ClusterFirst"</b> indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"ClusterFirstWithHostNet"</b> indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"Default"</b> indicates that the pod should use the default (as determined by kubelet) DNS settings. - <b>"None"</b> indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig.</p>
<b>enableServiceLinks</b>	<b>boolean</b>	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.

Property	Type	Description
<b>ephemeralContainers</b>	<b>array (EphemeralContainer)</b>	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's <code>ephemeralcontainers</code> subresource.
<b>hostAliases</b>	<b>array (HostAlias)</b>	<code>HostAliases</code> is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non- <code>hostNetwork</code> pods.
<b>hostIPC</b>	<b>boolean</b>	Use the host's ipc namespace. Optional: Default to false.
<b>hostNetwork</b>	<b>boolean</b>	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
<b>hostPID</b>	<b>boolean</b>	Use the host's pid namespace. Optional: Default to false.
<b>hostUsers</b>	<b>boolean</b>	Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with <code>CAP_SYS_MODULE</code> . When set to false, a new users is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the <code>UserNamespacesSupport</code> feature.

Property	Type	Description
<b>hostname</b>	<b>string</b>	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
<b>imagePullSecrets</b>	<b>array (LocalObjectReference)</b>	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod">https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod</a>
<b>initContainers</b>	<b>array (Container)</b>	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/init-containers/">https://kubernetes.io/docs/concepts/workloads/pods/init-containers/</a>

Property	Type	Description
<b>nodeName</b>	<b>string</b>	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
<b>nodeSelector</b>	<b>object (string)</b>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a>

Property	Type	Description
<b>os</b>	<b>PodOS</b>	<p>Specifies the OS of the containers in the pod. Some pod and container fields are restricted if this is set.</p> <p>If the OS field is set to linux, the following fields must be unset: - securityContext.windowsOptions</p> <p>If the OS field is set to windows, following fields must be unset: - spec.hostPID - spec.hostIPC - spec.hostUsers - spec.securityContext.seLinuxOptions - spec.securityContext.seccompProfile - spec.securityContext.fsGroup - spec.securityContext.fsGroupChangePolicy - spec.securityContext.sysctls - spec.shareProcessNamespace - spec.securityContext.runAsUser - spec.securityContext.runAsGroup - spec.securityContext.supplementalGroups - spec.containers[].<b>securityContext.seLinuxOptions</b> - <b>spec.containers[].securityContext.seccompProfile</b> - spec.containers[].<b>securityContext.capabilities</b> - <b>spec.containers[].securityContext.readOnlyRootFilesystem</b> - spec.containers[].<b>securityContext.privileged</b> - <b>spec.containers[].securityContext.allowPrivilegeEscalation</b> - spec.containers[].<b>securityContext.procMount</b> - <b>spec.containers[].securityContext.runAsUser</b> - spec.containers[*].securityContext.runAsGroup</p>

Property	Type	Description
<b>overhead</b>	<b>object (Quantity)</b>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md">https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md</a>
<b>preemptionPolicy</b>	<b>string</b>	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.  Possible enum values: - <b>"Never"</b> means that pod never preempts other pods with lower priority. - <b>"PreemptLowerPriority"</b> means that pod can preempt other pods with lower priority.
<b>priority</b>	<b>integer</b>	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

Property	Type	Description
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
<b>readinessGates</b>	<b>array (PodReadinessGate)</b>	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: <a href="https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates">https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates</a>
<b>resourceClaims</b>	<b>array (PodResourceClaim)</b>	ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable.



Property	Type	Description
<b>restartPolicy</b>	<b>string</b>	<p>Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy</a></p> <p>Possible enum values: - <b>"Always"</b> - <b>"Never"</b> - <b>"OnFailure"</b></p>
<b>runtimeClassName</b>	<b>string</b>	<p>RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class">https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class</a></p>
<b>schedulerName</b>	<b>string</b>	<p>If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.</p>
<b>schedulingGates</b>	<b>array (PodSchedulingGate)</b>	<p>SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.</p> <p>SchedulingGates can only be set at pod creation time, and be removed only afterwards.</p> <p>This is a beta feature enabled by the PodSchedulingReadiness feature gate.</p>

Property	Type	Description
<b>securityContext</b>	<b>PodSecurityContext</b>	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
<b>serviceAccount</b>	<b>string</b>	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
<b>serviceAccountName</b>	<b>string</b>	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: <a href="https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/">https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/</a>
<b>setHostnameAsFQDN</b>	<b>boolean</b>	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
<b>shareProcessNamespace</b>	<b>boolean</b>	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.

Property	Type	Description
<b>subdomain</b>	<b>string</b>	If specified, the fully qualified Pod hostname will be "<hostname>.<subdomain>.<pod namespace>.svc.<cluster domain>". If not specified, the pod will not have a domainname at all.
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
<b>tolerations</b>	<b>array (Toleration)</b>	If specified, the pod's tolerations.
<b>topologySpreadConstraints</b>	<b>array (TopologySpreadConstraint)</b>	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
<b>volumes</b>	<b>array (Volume)</b>	List of volumes that can be mounted by containers belonging to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes">https://kubernetes.io/docs/concepts/storage/volumes</a>

#### 4.1.98. io.k8s.api.core.v1.PodTemplateList schema

##### Description

PodTemplateList is a list of PodTemplates.

##### Type

**object**

**Required**

- **items**

**Schema**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (PodTemplate)</b>	List of pod templates
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

**4.1.99. io.k8s.api.core.v1.PodTemplateSpec schema****Description**

PodTemplateSpec describes the data a pod should have when created from a template

**Type**

**object**

**Schema**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>metadata</b>	<a href="#">ObjectMeta</a>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<a href="#">PodSpec</a>	Specification of the desired behavior of the pod. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status</a>

#### 4.1.100. io.k8s.api.core.v1.PortworxVolumeSource schema

##### Description

PortworxVolumeSource represents a Portworx volume resource.

##### Type

**object**

##### Required

- **volumelD**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>volumelD</b>	<b>string</b>	volumelD uniquely identifies a Portworx volume

#### 4.1.101. io.k8s.api.core.v1.PreferredSchedulingTerm schema

## Description

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

## Type

**object**

## Required

- **weight**
- **preference**

## Schema

Property	Type	Description
<b>preference</b>	<a href="#">NodeSelectorTerm</a>	A node selector term, associated with the corresponding weight.
<b>weight</b>	<b>integer</b>	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

## 4.1.102. io.k8s.api.core.v1.Probe schema

### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

### Type

**object**

### Schema

Property	Type	Description
<b>exec</b>	<a href="#">ExecAction</a>	Exec specifies the action to take.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<a href="#">GRPCAction</a>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<a href="#">HTTPGetAction</a>	HTTPGet specifies the http request to perform.

Property	Type	Description
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>TCPSocketAction</b>	TCPSocket specifies an action involving a TCP port.
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.

Property	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 4.1.103. io.k8s.api.core.v1.ProjectedVolumeSource schema

##### Description

Represents a projected volume source

##### Type

**object**

##### Schema

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>sources</b>	<b>array (VolumeProjection)</b>	sources is the list of volume projections

#### 4.1.104. io.k8s.api.core.v1.QuobyteVolumeSource schema

##### Description

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **registry**



- **volume**

#### Schema

Property	Type	Description
<b>group</b>	<b>string</b>	group to map volume access to Default is no group
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
<b>registry</b>	<b>string</b>	registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
<b>tenant</b>	<b>string</b>	tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
<b>user</b>	<b>string</b>	user to map volume access to Defaults to serviceaccount user
<b>volume</b>	<b>string</b>	volume is a string that references an already created Quobyte volume by name.

#### 4.1.105. io.k8s.api.core.v1.RBDPersistentVolumeSource schema

##### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**
- **image**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#rbd">https://kubernetes.io/docs/concepts/storage/volumes#rbd</a>
<b>image</b>	<b>string</b>	image is the rados image name. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>keyring</b>	<b>string</b>	keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>monitors</b>	<b>array (string)</b>	monitors is a collection of Ceph monitors. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>pool</b>	<b>string</b>	pool is the rados pool name. Default is rbd. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>SecretReference</b>	secretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>user</b>	<b>string</b>	user is the rados user name. Default is admin. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

## 4.1.106. io.k8s.api.core.v1.RBDVolumeSource schema

### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

### Type

**object**

### Required

- **monitors**
- **image**

### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#rbd">https://kubernetes.io/docs/concepts/storage/volumes#rbd</a>
<b>image</b>	<b>string</b>	image is the rados image name. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>keyring</b>	<b>string</b>	keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>monitors</b>	<b>array (string)</b>	monitors is a collection of Ceph monitors. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>pool</b>	<b>string</b>	pool is the rados pool name. Default is rbd. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>LocalObjectReference</b>	secretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>user</b>	<b>string</b>	user is the rados user name. Default is admin. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

#### 4.1.107. io.k8s.api.core.v1.ReplicationControllerList schema

##### Description

ReplicationControllerList is a collection of replication controllers.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>items</b>	<b>array (ReplicationController)</b>	List of replication controllers. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.108. io.k8s.api.core.v1.ResourceClaim schema

##### Description

ResourceClaim references one entry in PodSpec.ResourceClaims.

##### Type

**object**

##### Required

- **name**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container.

#### 4.1.109. io.k8s.api.core.v1.ResourceFieldSelector schema

**Description**

ResourceFieldSelector represents container resources (cpu, memory) and their output format

**Type**

**object**

**Required**

- **resource**

**Schema**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

**4.1.110. io.k8s.api.core.v1.ResourceQuotaList schema****Description**

ResourceQuotaList is a list of ResourceQuota items.

**Type**

**object**

**Required**

- **items**

**Schema**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>items</b>	<b>array (ResourceQuota)</b>	Items is a list of ResourceQuota objects. More info: <a href="https://kubernetes.io/docs/concepts/policy/resource-quotas/">https://kubernetes.io/docs/concepts/policy/resource-quotas/</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.111. io.k8s.api.core.v1.ResourceRequirements schema

##### Description

ResourceRequirements describes the compute resource requirements.

##### Type

**object**

##### Schema

Property	Type	Description
<b>claims</b>	<b>array (ResourceClaim)</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>

Property	Type	Description
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 4.1.112. io.k8s.api.core.v1.ScaleIOPersistentVolumeSource schema

##### Description

ScaleIOPersistentVolumeSource represents a persistent ScaleIO volume

##### Type

**object**

##### Required

- **gateway**
- **system**
- **secretRef**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs"
<b>gateway</b>	<b>string</b>	gateway is the host address of the ScaleIO API Gateway.



Property	Type	Description
<b>protectionDomain</b>	<b>string</b>	protectionDomain is the name of the ScaleIO Protection Domain for the configured storage.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>SecretReference</b>	secretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
<b>sslEnabled</b>	<b>boolean</b>	sslEnabled is the flag to enable/disable SSL communication with Gateway, default false
<b>storageMode</b>	<b>string</b>	storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
<b>storagePool</b>	<b>string</b>	storagePool is the ScaleIO Storage Pool associated with the protection domain.
<b>system</b>	<b>string</b>	system is the name of the storage system as configured in ScaleIO.
<b>volumeName</b>	<b>string</b>	volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source.

#### 4.1.113. io.k8s.api.core.v1.ScaleIOVolumeSource schema

##### Description

ScaleIOVolumeSource represents a persistent ScaleIO volume

##### Type

**object**

##### Required

- **gateway**

- **system**
- **secretRef**

## Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
<b>gateway</b>	<b>string</b>	gateway is the host address of the ScaleIO API Gateway.
<b>protectionDomain</b>	<b>string</b>	protectionDomain is the name of the ScaleIO Protection Domain for the configured storage.
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>LocalObjectReference</b>	secretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
<b>sslEnabled</b>	<b>boolean</b>	sslEnabled Flag enable/disable SSL communication with Gateway, default false
<b>storageMode</b>	<b>string</b>	storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
<b>storagePool</b>	<b>string</b>	storagePool is the ScaleIO Storage Pool associated with the protection domain.
<b>system</b>	<b>string</b>	system is the name of the storage system as configured in ScaleIO.

Property	Type	Description
<b>volumeName</b>	<b>string</b>	volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source.

#### 4.1.114. io.k8s.api.core.v1.SeccompProfile schema

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

##### Schema

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".

Property	Type	Description
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 4.1.115. io.k8s.api.core.v1.SecretEnvSource schema

##### Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

##### Type

**object**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

#### 4.1.116. io.k8s.api.core.v1.SecretKeySelector schema

##### Description

SecretKeySelector selects a key of a Secret.

#### Type

**object**

#### Required

- **key**

#### Schema

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

### 4.1.117. io.k8s.api.core.v1.SecretList schema

#### Description

SecretList is a list of Secret.

#### Type

**object**

#### Required

- **items**

#### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>items</b>	<b>array (Secret)</b>	Items is a list of secret objects. More info: <a href="https://kubernetes.io/docs/concepts/configuration/secret">https://kubernetes.io/docs/concepts/configuration/secret</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.118. io.k8s.api.core.v1.SecretProjection schema

##### Description

Adapts a secret into a projected volume.

The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>items</b>	<b>array (KeyToPath)</b>	items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its key must be defined

#### 4.1.119. io.k8s.api.core.v1.SecretReference schema

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

##### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 4.1.120. io.k8s.api.core.v1.SecretVolumeSource schema

##### Description

Adapts a Secret into a volume.

The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

## Type

**object**

## Schema

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array (KeyToPath)</b>	items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its keys must be defined
<b>secretName</b>	<b>string</b>	secretName is the name of the secret in the pod's namespace to use. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#secret">https://kubernetes.io/docs/concepts/storage/volumes#secret</a>



## 4.1.121. io.k8s.api.core.v1.SecurityContext schema

### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

### Type

**object**

### Schema

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>Capabilities</b>	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime. Note that this field cannot be set when spec.os.name is windows.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>SELinuxOptions</b>	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seccompProfile</b>	<b>SeccompProfile</b>	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>windowsOptions</b>	<a href="#">WindowsSecurityContextOptions</a>	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is linux.

#### 4.1.122. io.k8s.api.core.v1.SELinuxOptions schema

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

##### Schema

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 4.1.123. io.k8s.api.core.v1.ServiceAccountList schema

##### Description

ServiceAccountList is a list of ServiceAccount objects

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ServiceAccount)</b>	List of ServiceAccounts. More info: <a href="https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/">https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.124. io.k8s.api.core.v1.ServiceAccountTokenProjection schema

##### Description

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

##### Type

**object**

##### Required

- **path**

##### Schema

Property	Type	Description
<b>audience</b>	<b>string</b>	audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
<b>expirationSeconds</b>	<b>integer</b>	expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
<b>path</b>	<b>string</b>	path is the path relative to the mount point of the file to project the token into.

#### 4.1.125. io.k8s.api.core.v1.ServiceList schema

##### Description

ServiceList holds a list of services.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Service)</b>	List of services
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.126. io.k8s.api.core.v1.StorageOSPersistentVolumeSource schema

##### Description

Represents a StorageOS persistent volume resource.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>ObjectReference</b>	secretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
<b>volumeName</b>	<b>string</b>	volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
<b>volumeNamespace</b>	<b>string</b>	volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

#### 4.1.127. io.k8s.api.core.v1.StorageOSVolumeSource schema

##### Description

Represents a StorageOS persistent volume resource.

##### Type

**object**

##### Schema



Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>LocalObjectReference</b>	secretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
<b>volumeName</b>	<b>string</b>	volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
<b>volumeNamespace</b>	<b>string</b>	volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

#### 4.1.128. io.k8s.api.core.v1.Sysctl schema

##### Description

Sysctl defines a kernel parameter to be set

##### Type

**object**

##### Required

- **name**

- **value**

#### Schema

Property	Type	Description
<b>name</b>	<b>string</b>	Name of a property to set
<b>value</b>	<b>string</b>	Value of a property to set

### 4.1.129. io.k8s.api.core.v1.TCPSocketAction schema

#### Description

TCPSocketAction describes an action based on opening a socket

#### Type

**object**

#### Required

- **port**

#### Schema

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

### 4.1.130. io.k8s.api.core.v1.Toleration schema

#### Description

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

#### Type

**object**

#### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>effect</b>	<b>string</b>	<p>Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.</p> <p>Possible enum values: -  <b>"NoExecute"</b> Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController. -  <b>"NoSchedule"</b> Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler. -  <b>"PreferNoSchedule"</b> Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler.</p>
<b>key</b>	<b>string</b>	<p>Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.</p>
<b>operator</b>	<b>string</b>	<p>Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.</p> <p>Possible enum values: - <b>"Equal"</b>  - <b>"Exists"</b></p>

Property	Type	Description
<b>tolerationSeconds</b>	<b>integer</b>	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
<b>value</b>	<b>string</b>	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

#### 4.1.131. io.k8s.api.core.v1.TopologySelectorLabelRequirement schema

##### Description

A topology selector requirement is a selector that matches given label. This is an alpha feature and may change in the future.

##### Type

**object**

##### Required

- **key**
- **values**

##### Schema

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>values</b>	<b>array (string)</b>	An array of string values. One value must match the label to be selected. Each entry in Values is ORed.

#### 4.1.132. io.k8s.api.core.v1.TopologySelectorTerm schema

##### Description

A topology selector term represents the result of label queries. A null or empty topology selector term matches no objects. The requirements of them are ANDed. It provides a subset of functionality as NodeSelectorTerm. This is an alpha feature and may change in the future.

#### Type

**object**

#### Schema

Property	Type	Description
<b>matchLabelExpressions</b>	<b>array</b> ( <b>TopologySelectorLabelRequirement</b> )	A list of topology selector requirements by labels.

### 4.1.133. io.k8s.api.core.v1.TopologySpreadConstraint schema

#### Description

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

#### Type

**object**

#### Required

- **maxSkew**
- **topologyKey**
- **whenUnsatisfiable**

#### Schema

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

Property	Type	Description
<b>matchLabelKeys</b>	<b>array (string)</b>	<p>MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.</p> <p>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).</p>

Property	Type	Description
<b>maxSkew</b>	<b>integer</b>	<p>MaxSkew describes the degree to which pods may be unevenly distributed. When <b>whenUnsatisfiable=DoNotSchedule</b>, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1.   zone1   zone2   zone3     P P   P P   P   - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <b>whenUnsatisfiable=ScheduleAnyway</b>, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.</p>

Property	Type	Description
<b>minDomains</b>	<b>integer</b>	<p>MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.</p> <p>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2:   zone1   zone2   zone3     P P   P P   P P   The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew.</p> <p>This is a beta field and requires the MinDomainsInPodTopologySpread feature gate to be enabled (enabled by default).</p>



Property	Type	Description
<b>nodeAffinityPolicy</b>	<b>string</b>	<p>NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.</p> <p>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>
<b>nodeTaintsPolicy</b>	<b>string</b>	<p>NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.</p> <p>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>

Property	Type	Description
<b>topologyKey</b>	<b>string</b>	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field.

Property	Type	Description
<b>whenUnsatisfiable</b>	<b>string</b>	<p>WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1:   zone1   zone2   zone3     P P P   P   P   If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <b>more</b> imbalanced. It's a required field.</p> <p>Possible enum values: - <b>"DoNotSchedule"</b> instructs the scheduler not to schedule the pod when constraints are not satisfied. - <b>"ScheduleAnyway"</b> instructs the scheduler to schedule the pod even if constraints are not satisfied.</p>

#### 4.1.134. io.k8s.api.core.v1.TypedLocalObjectReference schema

##### Description

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

##### Type

**object**

##### Required

- **kind**

- **name**

#### Schema

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

#### 4.1.135. io.k8s.api.core.v1.TypedObjectReference schema

##### Description

##### Type

**object**

##### Required

- **kind**
- **name**

#### Schema

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

#### 4.1.136. io.k8s.api.core.v1.Volume schema

##### Description

Volume represents a named volume in a pod that may be accessed by any container in the pod.

##### Type

**object**

##### Required

- **name**

##### Schema

Property	Type	Description
<b>awsElasticBlockStore</b>	<b><a href="#">AWSElasticBlockStoreVolumeSource</a></b>	awsElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>azureDisk</b>	<b><a href="#">AzureDiskVolumeSource</a></b>	azureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
<b>azureFile</b>	<b><a href="#">AzureFileVolumeSource</a></b>	azureFile represents an Azure File Service mount on the host and bind mount to the pod.

Property	Type	Description
<b>cephfs</b>	<b>CephFSVolumeSource</b>	cephFS represents a Ceph FS mount on the host that shares a pod's lifetime
<b>cinder</b>	<b>CinderVolumeSource</b>	cinder represents a cinder volume attached and mounted on kubelets host machine. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>configMap</b>	<b>ConfigMapVolumeSource</b>	configMap represents a configMap that should populate this volume
<b>csi</b>	<b>CSIVolumeSource</b>	csi (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
<b>downwardAPI</b>	<b>DownwardAPIVolumeSource</b>	downwardAPI represents downward API about the pod that should populate this volume
<b>emptyDir</b>	<b>EmptyDirVolumeSource</b>	emptyDir represents a temporary directory that shares a pod's lifetime. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

Property	Type	Description
<b>ephemeral</b>	<b><a href="#">EphemeralVolumeSource</a></b>	<p>ephemeral represents a volume that is handled by a cluster storage driver. The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see <a href="#">EphemeralVolumeSource</a> for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
<b>fc</b>	<b><a href="#">FCVolumeSource</a></b>	<p>fc represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>
<b>flexVolume</b>	<b><a href="#">FlexVolumeSource</a></b>	<p>flexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.</p>

Property	Type	Description
<b>flocker</b>	<b>FlockerVolumeSource</b>	flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
<b>gcePersistentDisk</b>	<b>GCEPersistentDiskVolumeSource</b>	gcePersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersisitentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersisitentdisk</a>
<b>gitRepo</b>	<b>GitRepoVolumeSource</b>	gitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
<b>glusterfs</b>	<b>GlusterfsVolumeSource</b>	glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md">https://examples.k8s.io/volumes/glusterfs/README.md</a>
<b>hostPath</b>	<b>HostPathVolumeSource</b>	hostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>
<b>iscsi</b>	<b>ISCSIVolumeSource</b>	iscsi represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: <a href="https://examples.k8s.io/volumes/iscsi/README.md">https://examples.k8s.io/volumes/iscsi/README.md</a>



Property	Type	Description
<b>name</b>	<b>string</b>	name of the volume. Must be a DNS_LABEL and unique within the pod. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>nfs</b>	<b>NFSVolumeSource</b>	nfs represents an NFS mount on the host that shares a pod's lifetime More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>persistentVolumeClaim</b>	<b>PersistentVolumeClaimVolumeSource</b>	persistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims">https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims</a>
<b>photonPersistentDisk</b>	<b>PhotonPersistentDiskVolumeSource</b>	photonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
<b>portworxVolume</b>	<b>PortworxVolumeSource</b>	portworxVolume represents a portworx volume attached and mounted on kubelets host machine
<b>projected</b>	<b>ProjectedVolumeSource</b>	projected items for all in one resources secrets, configmaps, and downward API
<b>quobyte</b>	<b>QuobyteVolumeSource</b>	quobyte represents a Quobyte mount on the host that shares a pod's lifetime
<b>rbd</b>	<b>RBDVolumeSource</b>	rbd represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md">https://examples.k8s.io/volumes/rbd/README.md</a>
<b>scaleIO</b>	<b>ScaleIOVolumeSource</b>	scaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.

Property	Type	Description
<b>secret</b>	<a href="#">SecretVolumeSource</a>	secret represents a secret that should populate this volume. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#secret">https://kubernetes.io/docs/concepts/storage/volumes#secret</a>
<b>storageos</b>	<a href="#">StorageOSVolumeSource</a>	storageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
<b>vsphereVolume</b>	<a href="#">VsphereVirtualDiskVolumeSource</a>	vsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

#### 4.1.137. io.k8s.api.core.v1.VolumeDevice schema

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

##### Required

- **name**
- **devicePath**

##### Schema

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

#### 4.1.138. io.k8s.api.core.v1.VolumeMount schema

##### Description

VolumeMount describes a mounting of a Volume within a container.

##### Type

**object**

##### Required

- **name**

- `mountPath`

## Schema

Property	Type	Description
<code>mountPath</code>	<code>string</code>	Path within the container at which the volume should be mounted. Must not contain ':'.
<code>mountPropagation</code>	<code>string</code>	<p><code>mountPropagation</code> determines how mounts are propagated from the host to container and the other way around. When not set, <code>MountPropagationNone</code> is used. This field is beta in 1.10.</p> <p>Possible enum values: - <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). - <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<code>name</code>	<code>string</code>	This must match the Name of a Volume.
<code>readOnly</code>	<code>boolean</code>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.

Property	Type	Description
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 4.1.139. io.k8s.api.core.v1.VolumeNodeAffinity schema

##### Description

VolumeNodeAffinity defines constraints that limit what nodes this volume can be accessed from.

##### Type

**object**

##### Schema

Property	Type	Description
<b>required</b>	<a href="#">NodeSelector</a>	required specifies hard node constraints that must be met.

#### 4.1.140. io.k8s.api.core.v1.VolumeProjection schema

##### Description

Projection that may be projected along with other supported volume types

##### Type

**object**

##### Schema

Property	Type	Description
<b>configMap</b>	<a href="#">ConfigMapProjection</a>	configMap information about the configMap data to project
<b>downwardAPI</b>	<a href="#">DownwardAPIProjection</a>	downwardAPI information about the downwardAPI data to project

Property	Type	Description
<b>secret</b>	<a href="#">SecretProjection</a>	secret information about the secret data to project
<b>serviceAccountToken</b>	<a href="#">ServiceAccountTokenProjection</a>	serviceAccountToken is information about the serviceAccountToken data to project

#### 4.1.141. io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource schema

##### Description

Represents a vSphere volume resource.

##### Type

**object**

##### Required

- **volumePath**

##### Schema

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>storagePolicyID</b>	<b>string</b>	storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
<b>storagePolicyName</b>	<b>string</b>	storagePolicyName is the storage Policy Based Management (SPBM) profile name.
<b>volumePath</b>	<b>string</b>	volumePath is the path that identifies vSphere volume vmdk

#### 4.1.142. io.k8s.api.core.v1.WeightedPodAffinityTerm schema

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

## Type

**object**

## Required

- **weight**
- **podAffinityTerm**

## Schema

Property	Type	Description
<b>podAffinityTerm</b>	<b>PodAffinityTerm</b>	Required. A pod affinity term, associated with the corresponding weight.
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

## 4.1.143. io.k8s.api.core.v1.WindowsSecurityContextOptions schema

## Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

## Type

**object**

## Schema

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.

Property	Type	Description
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 4.1.144. io.k8s.api.discovery.v1.EndpointSliceList schema

##### Description

EndpointSliceList represents a list of endpoint slices

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (EndpointSlice)</b>	items is the list of endpoint slices
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata.

#### 4.1.145. io.k8s.api.events.v1.EventList schema

##### Description

EventList is a list of Event objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Event)</b>	items is a list of schema objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.146. io.k8s.api.flowcontrol.v1beta3.FlowSchemaList schema

##### Description

FlowSchemaList is a list of FlowSchema objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (FlowSchema)</b>	<b>items</b> is a list of FlowSchemas.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	<b>metadata</b> is the standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.147. io.k8s.api.flowcontrol.v1beta3.PriorityLevelConfigurationList schema

##### Description

PriorityLevelConfigurationList is a list of PriorityLevelConfiguration objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (PriorityLevelConfiguration)</b>	<b>items</b> is a list of request-priorities.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	<b>metadata</b> is the standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.148. io.k8s.api.networking.v1.IngressClassList schema

##### Description

IngressClassList is a collection of IngressClasses.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (IngressClass)</b>	items is the list of IngressClasses.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata.

#### 4.1.149. io.k8s.api.networking.v1.IngressList schema

##### Description

IngressList is a collection of Ingress.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Ingress)</b>	items is the list of Ingress.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.150. io.k8s.api.networking.v1.NetworkPolicyList schema

##### Description

NetworkPolicyList is a list of NetworkPolicy objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (NetworkPolicy)</b>	items is a list of schema objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.151. io.k8s.api.node.v1.RuntimeClassList schema

##### Description

RuntimeClassList is a list of RuntimeClass objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (RuntimeClass)</b>	items is a list of schema objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.152. io.k8s.api.policy.v1.PodDisruptionBudgetList schema

##### Description

PodDisruptionBudgetList is a collection of PodDisruptionBudgets.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (PodDisruptionBudget)</b>	Items is a list of PodDisruptionBudgets
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.153. io.k8s.api.rbac.v1.ClusterRoleBindingList schema

##### Description

ClusterRoleBindingList is a collection of ClusterRoleBindings

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ClusterRoleBinding)</b>	Items is a list of ClusterRoleBindings
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata.

#### 4.1.154. io.k8s.api.rbac.v1.ClusterRoleList schema

##### Description

ClusterRoleList is a collection of ClusterRoles

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (ClusterRole)</b>	Items is a list of ClusterRoles
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata.

#### 4.1.155. io.k8s.api.rbac.v1.RoleBindingList schema

##### Description

RoleBindingList is a collection of RoleBindings

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (RoleBinding)</b>	Items is a list of RoleBindings
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata.

#### 4.1.156. io.k8s.api.rbac.v1.RoleList schema

##### Description

RoleList is a collection of Roles

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Role)</b>	Items is a list of Roles
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata.

#### 4.1.157. io.k8s.api.scheduling.v1.PriorityClassList schema

##### Description

PriorityClassList is a collection of priority classes.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (PriorityClass)</b>	items is the list of PriorityClasses
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.158. io.k8s.api.storage.v1.CSIDriverList schema

##### Description

CSIDriverList is a collection of CSIDriver objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (CSIDriver)</b>	items is the list of CSIDriver
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.159. io.k8s.api.storage.v1.CSINodeList schema

##### Description

CSINodeList is a collection of CSINode objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (CSINode)</b>	items is the list of CSINode
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.160. io.k8s.api.storage.v1.CSISStorageCapacityList schema

##### Description

CSISStorageCapacityList is a collection of CSISStorageCapacity objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (CSIStorageCapacity)</b>	items is the list of CSIStorageCapacity objects.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.161. io.k8s.api.storage.v1.StorageClassList schema

##### Description

StorageClassList is a collection of storage classes.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (StorageClass)</b>	items is the list of StorageClasses
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.162. io.k8s.api.storage.v1.VolumeAttachmentList schema

##### Description

VolumeAttachmentList is a collection of VolumeAttachment objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (VolumeAttachment)</b>	items is the list of VolumeAttachments
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.163. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.CustomResourceDefinitionList schema

##### Description

CustomResourceDefinitionList is a list of CustomResourceDefinition objects.

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (CustomResourceDefinition)</b>	items list individual CustomResourceDefinition objects
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard object's metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.164. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.ExternalDocumentation schema

##### Description

ExternalDocumentation allows referencing an external resource for extended documentation.

##### Type

**object**

##### Schema

Property	Type	Description
<b>description</b>	<b>string</b>	

Property	Type	Description
<b>url</b>	<b>string</b>	

#### 4.1.165. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSON schema

##### Description

JSON represents any valid JSON value. These types are supported: bool, int64, float64, string, []interface{}, map[string]interface{} and nil.

##### Type

..

#### 4.1.166. io.k8s.apixtensions-apiserver.pkg.apis.apixtensions.v1.JSONSchemaProps schema

##### Description

JSONSchemaProps is a JSON-Schema following Specification Draft 4 (<http://json-schema.org/>).

##### Type

**object**

##### Schema

Property	Type	Description
<b>\$ref</b>	<b>string</b>	
<b>\$schema</b>	<b>string</b>	
<b>additionalItems</b>	<b>JSONSchemaPropsOrBool</b>	
<b>additionalProperties</b>	<b>JSONSchemaPropsOrBool</b>	
<b>allOf</b>	<b>array (undefined)</b>	
<b>anyOf</b>	<b>array (undefined)</b>	
<b>default</b>	<b>JSON</b>	default is a default value for undefined object fields. Defaulting is a beta feature under the CustomResourceDefaulting feature gate. Defaulting requires spec.preserveUnknownFields to be false.
<b>definitions</b>	<b>object (undefined)</b>	
<b>dependencies</b>	<b>object (undefined)</b>	

Property	Type	Description
<b>description</b>	<b>string</b>	
<b>enum</b>	<b>array (JSON)</b>	
<b>example</b>	<b>JSON</b>	
<b>exclusiveMaximum</b>	<b>boolean</b>	
<b>exclusiveMinimum</b>	<b>boolean</b>	
<b>externalDocs</b>	<b>ExternalDocumentation</b>	
<b>format</b>	<b>string</b>	<p>format is an OpenAPI v3 format string. Unknown formats are ignored. The following formats are validated:</p> <ul style="list-style-type: none"> <li>- bsonobjectid: a bson object ID, i.e. a 24 characters hex string</li> <li>- uri: an URI as parsed by Golang <code>net/url.ParseRequestURI</code></li> <li>- email: an email address as parsed by Golang <code>net/mail.ParseAddress</code></li> <li>- hostname: a valid representation for an Internet host name, as defined by RFC 1034, section 3.1 [RFC1034].</li> <li>- ipv4: an IPv4 IP as parsed by Golang <code>net.ParseIP</code></li> <li>- ipv6: an IPv6 IP as parsed by Golang <code>net.ParseIP</code></li> <li>- cidr: a CIDR as parsed by Golang <code>net.ParseCIDR</code></li> <li>- mac: a MAC address as parsed by Golang <code>net.ParseMAC</code></li> <li>- uuid: an UUID that allows uppercase defined by the regex <code>(?i)^[0-9a-f]{8}-?[0-9a-f]{4}-?[0-9a-f]{4}-?[0-9a-f]{4}-?[0-9a-f]{12}\$</code></li> <li>- uuid3: an UUID3 that allows uppercase defined by the regex <code>(?i)^[0-9a-f]{8}-?[0-9a-f]{4}-?3[0-9a-f]{3}-?[0-9a-f]{4}-?[0-9a-f]{12}\$</code></li> <li>- uuid4: an UUID4 that allows uppercase defined by the regex <code>(?i)^[0-9a-f]{8}-?[0-9a-f]{4}-?4[0-9a-f]{3}-?[89ab][0-9a-f]{3}-?[0-9a-f]{12}\$</code></li> <li>- uuid5: an UUID5 that allows uppercase defined by the regex <code>(?i)^[0-9a-f]{8}-?[0-9a-f]{4}-?5[0-9a-f]{3}-?[89ab][0-9a-f]{3}-?[0-9a-f]{12}\$</code></li> </ul>

Property	Type	Description
		<p>- isbn: an ISBN10 or ISBN13 number string like "0321751043" or "978-0321751041" - isbn10: an ISBN10 number string like "0321751043" - isbn13: an ISBN13 number string like "978-0321751041" - creditcard: a credit card number defined by the regex <code>^(?:4[0-9]{12}(?:[0-9]{3})?[5[1-5][0-9]{14} 6(?:011 5[0-9][0-9])[0-9]{12} 3[47][0-9]{13} 3(?:0[0-5]  [68][0-9])[0-9]{11} (?:2131 1800 35\d{3})\d{11})\$</code> with any non digit characters mixed in - ssn: a U.S. social security number following the regex <code>^\d{3}[- ]?\d{2}[- ]?\d{4}\$</code> - hexcolor: an hexadecimal color code like "FFFFFF: following the regex <code>^[0-9a-fA-F]{3}([0-9a-fA-F]{6})\$</code> - rgbcolor: an RGB color code like rgb like "rgb(255,255,255)" - byte: base64 encoded binary data - password: any kind of string - date: a date string like "2006-01-02" as defined by full-date in RFC3339 - duration: a duration string like "22 ns" as parsed by Golang time.ParseDuration or compatible with Scala duration format - datetime: a date time string like "2014-12-15T19:30:20.000Z" as defined by date-time in RFC3339.</p>
<b>id</b>	<b>string</b>	
<b>items</b>	<b><a href="#">JSONSchemaPropsOrArray</a></b>	
<b>maxItems</b>	<b>integer</b>	
<b>maxLength</b>	<b>integer</b>	
<b>maxProperties</b>	<b>integer</b>	
<b>maximum</b>	<b>number</b>	
<b>minItems</b>	<b>integer</b>	
<b>minLength</b>	<b>integer</b>	

Property	Type	Description
<b>minProperties</b>	<b>integer</b>	
<b>minimum</b>	<b>number</b>	
<b>multipleOf</b>	<b>number</b>	
<b>not</b>	<b>JSONSchemaProps</b>	
<b>nullable</b>	<b>boolean</b>	
<b>oneOf</b>	<b>array (undefined)</b>	
<b>pattern</b>	<b>string</b>	
<b>patternProperties</b>	<b>object (undefined)</b>	
<b>properties</b>	<b>object (undefined)</b>	
<b>required</b>	<b>array (string)</b>	
<b>title</b>	<b>string</b>	
<b>type</b>	<b>string</b>	
<b>uniqueItems</b>	<b>boolean</b>	
<b>x-kubernetes-embedded-resource</b>	<b>boolean</b>	x-kubernetes-embedded-resource defines that the value is an embedded Kubernetes runtime.Object, with TypeMeta and ObjectMeta. The type must be object. It is allowed to further restrict the embedded object. kind, apiVersion and metadata are validated automatically. x-kubernetes-preserve-unknown-fields is allowed to be true, but does not have to be if the object is fully specified (up to kind, apiVersion, metadata).

Property	Type	Description
<b>x-kubernetes-int-or-string</b>	<b>boolean</b>	<p>x-kubernetes-int-or-string specifies that this value is either an integer or a string. If this is true, an empty type is allowed and type as child of anyOf is permitted if following one of the following patterns:</p> <p>1) anyOf: - type: integer - type: string 2) allOf: - anyOf: - type: integer - type: string - ... zero or more</p>
<b>x-kubernetes-list-map-keys</b>	<b>array (string)</b>	<p>x-kubernetes-list-map-keys annotates an array with the x-kubernetes-list-type <b>map</b> by specifying the keys used as the index of the map.</p> <p>This tag <b>MUST</b> only be used on lists that have the "x-kubernetes-list-type" extension set to "map". Also, the values specified for this attribute must be a scalar typed field of the child structure (no nesting is supported).</p> <p>The properties specified must either be required or have a default value, to ensure those properties are present for all list items.</p>



Property	Type	Description
<b>x-kubernetes-list-type</b>	<b>string</b>	<p>x-kubernetes-list-type annotates an array to further describe its topology. This extension must only be used on lists and may have 3 possible values:</p> <p>1) <b>atomic</b>: the list is treated as a single entity, like a scalar. Atomic lists will be entirely replaced when updated. This extension may be used on any type of list (struct, scalar, ...). 2) <b>set</b>: Sets are lists that must not have multiple items with the same value. Each value must be a scalar, an object with x-kubernetes-map-type <b>atomic</b> or an array with x-kubernetes-list-type <b>atomic</b>. 3) <b>map</b>: These lists are like maps in that their elements have a non-index key used to identify them. Order is preserved upon merge. The map tag must only be used on a list with elements of type object. Defaults to atomic for arrays.</p>
<b>x-kubernetes-map-type</b>	<b>string</b>	<p>x-kubernetes-map-type annotates an object to further describe its topology. This extension must only be used when type is object and may have 2 possible values:</p> <p>1) <b>granular</b>: These maps are actual maps (key-value pairs) and each fields are independent from each other (they can each be manipulated by separate actors). This is the default behaviour for all maps. 2) <b>atomic</b>: the list is treated as a single entity, like a scalar. Atomic maps will be entirely replaced when updated.</p>

Property	Type	Description
<b>x-kubernetes-preserve-unknown-fields</b>	<b>boolean</b>	x-kubernetes-preserve-unknown-fields stops the API server decoding step from pruning fields which are not specified in the validation schema. This affects fields recursively, but switches back to normal pruning behaviour if nested properties or additionalProperties are specified in the schema. This can either be true or undefined. False is forbidden.
<b>x-kubernetes-validations</b>	<b>array (ValidationRule)</b>	x-kubernetes-validations describes a list of validation rules written in the CEL expression language. This field is an alpha-level. Using this field requires the feature gate <b>CustomResourceValidationExpressions</b> to be enabled.

#### 4.1.167. io.k8s.apimachinery.pkg.apis.extensions.v1.JSONSchemaPropsOrArray schema

##### Description

JSONSchemaPropsOrArray represents a value that can either be a JSONSchemaProps or an array of JSONSchemaProps. Mainly here for serialization purposes.

##### Type

..

#### 4.1.168. io.k8s.apimachinery.pkg.apis.extensions.v1.JSONSchemaPropsOrBool schema

##### Description

JSONSchemaPropsOrBool represents JSONSchemaProps or a boolean value. Defaults to true for the boolean property.

##### Type

..

#### 4.1.169. io.k8s.apimachinery.pkg.apis.extensions.v1.JSONSchemaPropsOrStringArray schema

##### Description

JSONSchemaPropsOrStringArray represents a JSONSchemaProps or a string array.

##### Type

..

## 4.1.170. io.k8s.apiextensions-apiserver.pkg.apis.apiextensions.v1.ValidationRule schema

### Description

ValidationRule describes a validation rule written in the CEL expression language.

### Type

**object**

### Required

- **rule**

### Schema

Property	Type	Description
<b>message</b>	<b>string</b>	Message represents the message displayed when validation fails. The message is required if the Rule contains line breaks. The message must not contain line breaks. If unset, the message is "failed rule: {Rule}". e.g. "must be a URL with the host matching spec.host"

Property	Type	Description
<b>messageExpression</b>	<b>string</b>	<p>MessageExpression declares a CEL expression that evaluates to the validation failure message that is returned when this rule fails. Since messageExpression is used as a failure message, it must evaluate to a string. If both message and messageExpression are present on a rule, then messageExpression will be used if validation fails. If messageExpression results in a runtime error, the runtime error is logged, and the validation failure message is produced as if the messageExpression field were unset. If messageExpression evaluates to an empty string, a string with only spaces, or a string that contains line breaks, then the validation failure message will also be produced as if the messageExpression field were unset, and the fact that messageExpression produced an empty string/string with only spaces/string with line breaks will be logged. messageExpression has access to all the same variables as the rule; the only difference is the return type. Example: "x must be less than max ("string(self.max)")"</p>
<b>rule</b>	<b>string</b>	<p>Rule represents the expression which will be evaluated by CEL. ref: <a href="https://github.com/google/cel-spec">https://github.com/google/cel-spec</a> The Rule is scoped to the location of the x-kubernetes-validations extension in the schema. The <b>self</b> variable in the CEL expression is bound to the scoped value. Example: - Rule scoped to the root of a resource with a status subresource: {"rule": "self.status.actual &lt;= self.spec.maxDesired"}</p> <p>If the Rule is scoped to an object with properties, the accessible properties of the object are field selectable via <b>self.field</b> and field</p>

Property	Type	Description
		<p>presence can be checked via <code>has(self.field)</code>. Null valued fields are treated as absent fields in CEL expressions. If the Rule is scoped to an object with <code>additionalProperties</code> (i.e. a map) the value of the map are accessible via <code>self[mapKey]</code>, map containment can be checked via <code>mapKey in self</code> and all entries of the map are accessible via CEL macros and functions such as <code>self.all(...)</code>. If the Rule is scoped to an array, the elements of the array are accessible via <code>self[i]</code> and also by macros and functions. If the Rule is scoped to a scalar, <code>self</code> is bound to the scalar value. Examples: - Rule scoped to a map of objects: <pre>{ "rule": "self.components['Widget'].priority &lt; 10" }</pre> - Rule scoped to a list of integers: <pre>{ "rule": "self.values.all(value, value &gt;= 0 &amp;&amp; value &lt; 100)" }</pre> - Rule scoped to a string value: <pre>{ "rule": "self.startsWith('kube')" }</pre></p> <p>The <code>apiVersion</code>, <code>kind</code>, <code>metadata.name</code> and <code>metadata.generateName</code> are always accessible from the root of the object and from any <code>x-kubernetes-embedded-resource</code> annotated objects. No other metadata properties are accessible.</p> <p>Unknown data preserved in custom resources via <code>x-kubernetes-preserve-unknown-fields</code> is not accessible in CEL expressions. This includes: - Unknown field values that are preserved by object schemas with <code>x-kubernetes-preserve-unknown-fields</code>. - Object properties where the property schema is of an "unknown type". An "unknown type" is recursively defined as: - A schema with no type and <code>x-kubernetes-preserve-unknown-fields</code> set to true - An array where the items schema is of an</p>

Property	Type	Description
		<p>"unknown type" - An object where the additionalProperties schema is of an "unknown type"</p> <p>Only property names of the form <b>[a-zA-Z_-]/[a-zA-Z0-9_-]/*</b> are accessible. Accessible property names are escaped according to the following rules when accessed in the expression: - ' escapes to <code>'_underscores_'</code> - '.' escapes to <code>'_dot_'</code> - '-' escapes to <code>'_dash_'</code> - '/' escapes to <code>'_slash_'</code> -</p> <p>Property names that exactly match a CEL RESERVED keyword escape to <code>'_{keyword}_'</code>. The keywords are: "true", "false", "null", "in", "as", "break", "const", "continue", "else", "for", "function", "if", "import", "let", "loop", "package", "namespace", "return".</p> <p>Examples: - Rule accessing a property named "namespace":  <code>{ "rule": self.__namespace__ → O }</code> - Rule accessing a property named "x-prop": <code>{ "rule": self.x__dash__prop → O }</code> - Rule accessing a property named "redactd": <code>{ "rule": self.redact__underscores__d → O }</code></p> <p>Equality on arrays with x-kubernetes-list-type of 'set' or 'map' ignores element order, i.e. <code>[1, 2] == [2, 1]</code>. Concatenation on arrays with x-kubernetes-list-type use the semantics of the list type:</p> <ul style="list-style-type: none"> <li>- 'set': <b>X + Y</b> performs a union where the array positions of all elements in <b>X</b> are preserved and non-intersecting elements in <b>Y</b> are appended, retaining their partial order.</li> <li>- 'map': <b>X + Y</b> performs a merge where the array positions of all keys in <b>X</b> are preserved but the values are overwritten by values in <b>Y</b> when the key sets of <b>X</b> and <b>Y</b> intersect. Elements in <b>Y</b> with non-intersecting keys are appended, retaining their partial order.</li> </ul>

#### 4.1.171. io.k8s.apimachinery.pkg.api.resource.Quantity schema

## Description

Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to `String()` and `AsInt64()` accessors. The serialization format is:

```
<quantity> ::= <signedNumber><suffix>
```

(Note that `<suffix>` may be empty, from the `""` case in `<decimalSI>`.)

```
<digit> ::= 0 | 1 | ... | 9
<digits> ::= <digit> | <digit><digits>
<number> ::= <digits> | <digits>.<digits> |
<digits>.|.<digits>
<sign> ::= "+" | "-"
<signedNumber> ::= <number> | <sign><number>
<suffix> ::=
<binarySI> | <decimalExponent> | <decimalSI>
<binarySI> ::= Ki | Mi | Gi | Ti | Pi | Ei
```

(International System of units; See: <http://physics.nist.gov/cuu/Units/binary.html>)

```
<decimalSI> ::= m | "" | k | M | G | T | P | E
```

(Note that  $1024 = 1\text{Ki}$  but  $1000 = 1\text{k}$ ; I didn't choose the capitalization.)

```
<decimalExponent> ::= "e" <signedNumber> | "E" <signedNumber>
```

No matter which of the three exponent forms is used, no quantity may represent a number greater than  $2^{63}-1$  in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: `0.1m` will rounded up to `1m`.) This may be extended in the future if we require larger or smaller quantities.

When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.

Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:

- No precision is lost - No fractional digits will be emitted - The exponent (or suffix) is as large as possible.

The sign will be omitted unless the number is negative.

Examples:

- `1.5` will be serialized as `"1500m"` - `1.5Gi` will be serialized as `"1536Mi"`

Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.

Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)

This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.

## Type

**string**

## 4.1.172. io.k8s.apimachinery.pkg.apis.meta.v1.Condition schema

### Description

Condition contains details for one aspect of the current state of this API Resource.

### Type

**object**

### Required

- **type**
- **status**
- **lastTransitionTime**
- **reason**
- **message**

### Schema

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	lastTransitionTime is the last time the condition transitioned from one status to another. This should be when the underlying condition changed. If that is not known, then using the time when the API field changed is acceptable.
<b>message</b>	<b>string</b>	message is a human readable message indicating details about the transition. This may be an empty string.
<b>observedGeneration</b>	<b>integer</b>	observedGeneration represents the .metadata.generation that the condition was set based upon. For instance, if .metadata.generation is currently 12, but the .status.conditions[x].observedGeneration is 9, the condition is out of date with respect to the current state of the instance.



Property	Type	Description
<b>reason</b>	<b>string</b>	reason contains a programmatic identifier indicating the reason for the condition's last transition. Producers of specific condition types may define expected values and meanings for this field, and whether the values are considered a guaranteed API. The value should be a CamelCase string. This field may not be empty.
<b>status</b>	<b>string</b>	status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	type of condition in CamelCase or in foo.example.com/CamelCase.

#### 4.1.173. io.k8s.apimachinery.pkg.apis.meta.v1.DeleteOptions schema

##### Description

DeleteOptions may be provided when deleting an API object.

##### Type

**object**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>dryRun</b>	<b>array (string)</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Property	Type	Description
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>preconditions</b>	<b>Preconditions</b>	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.

Property	Type	Description
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

#### 4.1.174. io.k8s.apimachinery.pkg.apis.meta.v1.FieldsV1 schema

##### Description

FieldsV1 stores a set of fields in a data structure like a Trie, in JSON format.

Each key is either a '.' representing the field itself, and will always map to an empty set, or a string representing a sub-field or item. The string will follow one of these four formats: 'f:<name>', where <name> is the name of a field in a struct, or key in a map 'v:<value>', where <value> is the exact json formatted value of a list item 'i:<index>', where <index> is position of a item in a list 'k:<keys>', where <keys> is a map of a list item's key fields to their unique values. If a key maps to an empty Fields value, the field that key represents is part of the set.

The exact format is defined in sigs.k8s.io/structured-merge-diff

##### Type

**object**

#### 4.1.175. io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector schema

##### Description

A label selector is a label query over a set of resources. The result of matchLabels and matchExpressions are ANDed. An empty label selector matches all objects. A null label selector matches no objects.

##### Type

**object**

##### Schema

Property	Type	Description
<b>matchExpressions</b>	<b>array</b> <b>(LabelSelectorRequirement)</b>	matchExpressions is a list of label selector requirements. The requirements are ANDed.

Property	Type	Description
<b>matchLabels</b>	<b>object (string)</b>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

#### 4.1.176. io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement schema

##### Description

A label selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

##### Type

**object**

##### Required

- **key**
- **operator**

##### Schema

Property	Type	Description
<b>key</b>	<b>string</b>	key is the label key that the selector applies to.
<b>operator</b>	<b>string</b>	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
<b>values</b>	<b>array (string)</b>	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

#### 4.1.177. io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta schema

##### Description

ListMeta describes metadata that synthetic resources must have, including lists and various status objects. A resource may have only one of {ObjectMeta, ListMeta}.

## Type

**object**

## Schema

Property	Type	Description
<b>continue</b>	<b>string</b>	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
<b>remainingItemCount</b>	<b>integer</b>	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <b>estimating</b> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

Property	Type	Description
<b>resourceVersion</b>	<b>string</b>	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>selfLink</b>	<b>string</b>	Deprecated: selfLink is a legacy read-only field that is no longer populated by the system.

#### 4.1.178. io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry schema

##### Description

ManagedFieldsEntry is a workflow-id, a FieldSet and the group version of the resource that the fieldset applies to.

##### Type

**object**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
<b>fieldsType</b>	<b>string</b>	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
<b>fieldsV1</b>	<b>FieldsV1</b>	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

Property	Type	Description
<b>manager</b>	<b>string</b>	Manager is an identifier of the workflow managing these fields.
<b>operation</b>	<b>string</b>	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
<b>subresource</b>	<b>string</b>	Subresource is the name of the subresource used to update that object, or empty string if the object was updated through the main resource. The value of this field is used to distinguish between managers, even if they share the same name. For example, a status update will be distinct from a regular update using the same manager name. Note that the APIVersion field is not related to the Subresource field and it always corresponds to the version of the main resource.
<b>time</b>	<b>Time</b>	Time is the timestamp of when the ManagedFields entry was added. The timestamp will also be updated if a field is added, the manager changes any of the owned fields value or removes a field. The timestamp does not update when a field is removed from the entry because another manager took it over.

#### 4.1.179. io.k8s.apimachinery.pkg.apis.meta.v1.MicroTime schema

##### Description

MicroTime is version of Time with microsecond level precision.

##### Type

**string**

#### 4.1.180. io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta schema

##### Description

ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.

##### Type

**object**

## Schema

Property	Type	Description
<b>annotations</b>	<b>object (string)</b>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/annotations">https://kubernetes.io/docs/concepts/overview/working-with-objects/annotations</a>
<b>creationTimestamp</b>	<b>Time</b>	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC.  Populated by the system. Read-only. Null for lists. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>deletionGracePeriodSeconds</b>	<b>integer</b>	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.



Property	Type	Description
<b>deletionTimestamp</b>	<b>Time</b>	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a></p>

Property	Type	Description
<b>finalizers</b>	<b>array (string)</b>	<p>Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.</p>

Property	Type	Description
<b>generateName</b>	<b>string</b>	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will return a 409.</p> <p>Applied only if Name is not specified. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</a></p>
<b>generation</b>	<b>integer</b>	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
<b>labels</b>	<b>object (string)</b>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels</a></p>

Property	Type	Description
<b>managedFields</b>	<b>array (ManagedFieldsEntry)</b>	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
<b>name</b>	<b>string</b>	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names#names</a>
<b>namespace</b>	<b>string</b>	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty.  Must be a DNS_LABEL. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces</a>

Property	Type	Description
<b>ownerReferences</b>	<b>array (OwnerReference)</b>	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
<b>resourceVersion</b>	<b>string</b>	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a></p>
<b>selfLink</b>	<b>string</b>	Deprecated: selfLink is a legacy read-only field that is no longer populated by the system.
<b>uid</b>	<b>string</b>	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids</a></p>

## 4.1.181. io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference schema

### Description

OwnerReference contains enough information to let you identify an owning object. An owning object must be in the same namespace as the dependent, or be cluster-scoped, so there is no namespace field.

### Type

**object**

### Required

- **apiVersion**
- **kind**
- **name**
- **uid**

### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>blockOwnerDeletion</b>	<b>boolean</b>	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. See <a href="https://kubernetes.io/docs/concepts/architecture/garbage-collection/#foreground-deletion">https://kubernetes.io/docs/concepts/architecture/garbage-collection/#foreground-deletion</a> for how the garbage collector interacts with this field and enforces the foreground deletion. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
<b>controller</b>	<b>boolean</b>	If true, this reference points to the managing controller.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names#names</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids</a>

#### 4.1.182. io.k8s.apimachinery.pkg.apis.meta.v1.Patch schema

##### Description

Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

##### Type

**object**

#### 4.1.183. io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions schema

##### Description

Preconditions must be fulfilled before an operation (update, delete, etc.) is carried out.

##### Type

**object**

##### Schema

Property	Type	Description
<b>resourceVersion</b>	<b>string</b>	Specifies the target ResourceVersion
<b>uid</b>	<b>string</b>	Specifies the target UID.

#### 4.1.184. io.k8s.apimachinery.pkg.apis.meta.v1.Status schema

##### Description

Status is a return value for calls that don't return other objects.

##### Type

**object**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>code</b>	<b>integer</b>	Suggested HTTP return code for this status, 0 if not set.
<b>details</b>	<b>StatusDetails</b>	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>message</b>	<b>string</b>	A human-readable description of the status of this operation.
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>reason</b>	<b>string</b>	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.



Property	Type	Description
<b>status</b>	<b>string</b>	Status of the operation. One of: "Success" or "Failure". More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status</a>

#### 4.1.185. io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause schema

##### Description

StatusCause provides more information about an api.Status failure, including cases when multiple errors are encountered.

##### Type

**object**

##### Schema

Property	Type	Description
<b>field</b>	<b>string</b>	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional.  Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
<b>message</b>	<b>string</b>	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
<b>reason</b>	<b>string</b>	A machine-readable description of the cause of the error. If this value is empty there is no information available.

#### 4.1.186. io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails schema

##### Description

StatusDetails is a set of additional properties that MAY be set by the server to provide additional information about a response. The Reason field of a Status object defines what attributes will be set.

Clients must ignore fields that do not match the defined type of each attribute, and should assume that any attribute may be empty, invalid, or under defined.

## Type

**object**

## Schema

Property	Type	Description
<b>causes</b>	<b>array (StatusCause)</b>	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
<b>group</b>	<b>string</b>	The group attribute of the resource associated with the status StatusReason.
<b>kind</b>	<b>string</b>	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
<b>retryAfterSeconds</b>	<b>integer</b>	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action - for those errors this field may indicate how long to wait before taking the alternate action.
<b>uid</b>	<b>string</b>	UID of the resource. (when there is a single resource which can be described). More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names#uids</a>

### 4.1.187. io.k8s.apimachinery.pkg.apis.meta.v1.Time schema

**Description**

Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.

**Type**

**string**

**4.1.188. io.k8s.apimachinery.pkg.apis.meta.v1.WatchEvent schema****Description**

Event represents a single event to a watched resource.

**Type**

**object**

**Required**

- **type**
- **object**

**Schema**

Property	Type	Description
<b>object</b>	<a href="#">RawExtension</a>	Object is: * If Type is Added or Modified: the new state of the object. * If Type is Deleted: the state of the object immediately before deletion. * If Type is Error: *Status is recommended; other types may make sense depending on context.
<b>type</b>	<b>string</b>	

**4.1.189. io.k8s.apimachinery.pkg.runtime.RawExtension schema****Description**

RawExtension is used to hold extensions in external versions.

To use this, make a field which has RawExtension as its type in your external, versioned struct, and Object in your internal struct. You also need to register your various plugin types.

```
type MyAPIObject struct {
    runtime.TypeMeta `json:",inline"`
    MyPlugin runtime.Object `json:"myPlugin"`
}
```

```
type PluginA struct {
    AOption string `json:"aOption"`
}
```

```
type MyAPIObject struct {
  runtime.TypeMeta `json:",inline"`
  MyPlugin runtime.RawExtension `json:"myPlugin"`
}
```

```
type PluginA struct {
  AOption string `json:"aOption"`
}
```

```
{
  "kind":"MyAPIObject",
  "apiVersion":"v1",
  "myPlugin": {
    "kind":"PluginA",
    "aOption":"foo",
  },
}
```

So what happens? Decode first uses json or yaml to unmarshal the serialized data into your external MyAPIObject. That causes the raw JSON to be stored, but not unpacked. The next step is to copy (using pkg/conversion) into the internal struct. The runtime package's DefaultScheme has conversion functions installed which will unpack the JSON stored in RawExtension, turning it into the correct object type, and storing it in the Object. (TODO: In the case where the object is of an unknown type, a runtime.Unknown object will be created and stored.)

## Type

**object**

### 4.1.190. io.k8s.apimachinery.pkg.util.intstr.IntOrString schema

#### Description

IntOrString is a type that can hold an int32 or a string. When used in JSON or YAML marshalling and unmarshalling, it produces or consumes the inner type. This allows you to have, for example, a JSON field that can accept a name or number.

## Type

**string**

### 4.1.191. io.k8s.kube-aggregator.pkg.apis.apiregistration.v1.APIServiceList schema

#### Description

APIServiceList is a list of APIService objects.

## Type

**object**

## Required

- **items**

## Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (APIService)</b>	Items is the list of APIService
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 4.1.192. io.k8s.migration.v1alpha1.StorageVersionMigrationList schema

##### Description

StorageVersionMigrationList is a list of StorageVersionMigration

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (StorageVersionMigration)</b>	List of storageversionmigrations. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.193. io.k8s.storage.snapshot.v1.VolumeSnapshotClassList schema

##### Description

VolumeSnapshotClassList is a list of VolumeSnapshotClass

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (VolumeSnapshotClass)</b>	List of volumesnapshotclasses. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.194. io.k8s.storage.snapshot.v1.VolumeSnapshotContentList schema

##### Description

VolumeSnapshotContentList is a list of VolumeSnapshotContent

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (VolumeSnapshotContent)</b>	List of volumesnapshotcontents. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.195. io.k8s.storage.snapshot.v1.VolumeSnapshotList schema

##### Description

VolumeSnapshotList is a list of VolumeSnapshot

##### Type

**object**

##### Required

- **items**

##### Schema



Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (VolumeSnapshot)</b>	List of volumesnapshots. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.196. io.openshift.internal.security.v1.RangeAllocationList schema

##### Description

RangeAllocationList is a list of RangeAllocation

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (RangeAllocation)</b>	List of rangeallocations. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.197. io.openshift.route.v1.RouteList schema

##### Description

RouteList is a list of Route

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (Route)</b>	List of routes. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.198. io.openshift.security.v1.SecurityContextConstraintsList schema

##### Description

SecurityContextConstraintsList is a list of SecurityContextConstraints

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (SecurityContextConstraints)</b>	List of securitycontextconstraints. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

#### 4.1.199. io.topolvm.v1.LogicalVolumeList schema

##### Description

LogicalVolumeList is a list of LogicalVolume

##### Type

**object**

##### Required

- **items**

##### Schema

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>items</b>	<b>array (LogicalVolume)</b>	List of logicalvolumes. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ListMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

## CHAPTER 5. EXTENSION APIS

### 5.1. API EXTENSIONS APIS

#### 5.1.1. CustomResourceDefinition [apiextensions.k8s.io/v1]

##### Description

CustomResourceDefinition represents a resource that should be exposed on the API server. Its name MUST be in the format <.spec.name>.<.spec.group>.

##### Type

**object**

### 5.2. CUSTOMRESOURCEDEFINITION [APIEXTENSIONS.K8S.IO/V1]

##### Description

CustomResourceDefinition represents a resource that should be exposed on the API server. Its name MUST be in the format <.spec.name>.<.spec.group>.

##### Type

**object**

##### Required

- **spec**

#### 5.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	CustomResourceDefinitionSpec describes how a user wants their resource to appear
<b>status</b>	<b>object</b>	CustomResourceDefinitionStatus indicates the state of the CustomResourceDefinition

### 5.2.1.1. .spec

#### Description

CustomResourceDefinitionSpec describes how a user wants their resource to appear

#### Type

**object**

#### Required

- **group**
- **names**
- **scope**
- **versions**

Property	Type	Description
<b>conversion</b>	<b>object</b>	CustomResourceConversion describes how to convert different versions of a CR.

Property	Type	Description
<b>group</b>	<b>string</b>	group is the API group of the defined custom resource. The custom resources are served under <b>/apis/&lt;group&gt;/...</b> . Must match the name of the CustomResourceDefinition (in the form <b>&lt;names.plural&gt;.&lt;group&gt;</b> ).
<b>names</b>	<b>object</b>	CustomResourceDefinitionNames indicates the names to serve this CustomResourceDefinition
<b>preserveUnknownFields</b>	<b>boolean</b>	preserveUnknownFields indicates that object fields which are not specified in the OpenAPI schema should be preserved when persisting to storage. apiVersion, kind, metadata and known fields inside metadata are always preserved. This field is deprecated in favor of setting <b>x-preserve-unknown-fields</b> to true in <b>spec.versions[*].schema.openAPIV3Schema</b> . See <a href="https://kubernetes.io/docs/tasks/extend-kubernetes/custom-resources/custom-resource-definitions/#field-pruning">https://kubernetes.io/docs/tasks/extend-kubernetes/custom-resources/custom-resource-definitions/#field-pruning</a> for details.
<b>scope</b>	<b>string</b>	scope indicates whether the defined custom resource is cluster- or namespace-scoped. Allowed values are <b>Cluster</b> and <b>Namespaced</b> .



Property	Type	Description
<b>versions</b>	<b>array</b>	versions is the list of all API versions of the defined custom resource. Version names are used to compute the order in which served versions are listed in API discovery. If the version string is "kube-like", it will sort above non "kube-like" version strings, which are ordered lexicographically. "Kube-like" versions start with a "v", then are followed by a number (the major version), then optionally the string "alpha" or "beta" and another number (the minor version). These are sorted first by GA > beta > alpha (where GA is a version with no suffix such as beta or alpha), and then by comparing major version, then minor version. An example sorted list of versions: v10, v2, v1, v11beta2, v10beta3, v3beta1, v12alpha1, v11alpha2, foo1, foo10.
<b>versions[]</b>	<b>object</b>	CustomResourceDefinitionVersion describes a version for CRD.

### 5.2.1.2. .spec.conversion

#### Description

CustomResourceConversion describes how to convert different versions of a CR.

#### Type

**object**

#### Required

- **strategy**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>strategy</b>	<b>string</b>	strategy specifies how custom resources are converted between versions. Allowed values are: - <b>"None"</b> : The converter only change the apiVersion and would not touch any other field in the custom resource. - <b>"Webhook"</b> : API Server will call to an external webhook to do the conversion. Additional information is needed for this option. This requires spec.preserveUnknownFields to be false, and spec.conversion.webhook to be set.
<b>webhook</b>	<b>object</b>	WebhookConversion describes how to call a conversion webhook

### 5.2.1.3. .spec.conversion.webhook

#### Description

WebhookConversion describes how to call a conversion webhook

#### Type

**object**

#### Required

- **conversionReviewVersions**

Property	Type	Description
<b>clientConfig</b>	<b>object</b>	WebhookClientConfig contains the information to make a TLS connection with the webhook.

Property	Type	Description
<b>conversionReviewVersions</b>	<b>array (string)</b>	conversionReviewVersions is an ordered list of preferred <b>ConversionReview</b> versions the Webhook expects. The API server will use the first version in the list which it supports. If none of the versions specified in this list are supported by API server, conversion will fail for the custom resource. If a persisted Webhook configuration specifies allowed versions and does not include any versions known to the API Server, calls to the webhook will fail.

#### 5.2.1.4. .spec.conversion.webhook.clientConfig

##### Description

WebhookClientConfig contains the information to make a TLS connection with the webhook.

##### Type

**object**

Property	Type	Description
<b>caBundle</b>	<b>string</b>	caBundle is a PEM encoded CA bundle which will be used to validate the webhook's server certificate. If unspecified, system trust roots on the apiserver are used.
<b>service</b>	<b>object</b>	ServiceReference holds a reference to Service.legacy.k8s.io

Property	Type	Description
<b>url</b>	<b>string</b>	<p>url gives the location of the webhook, in standard URL form (<b>scheme://host:port/path</b>). Exactly one of <b>url</b> or <b>service</b> must be specified.</p> <p>The <b>host</b> should not refer to a service running in the cluster; use the <b>service</b> field instead. The host might be resolved via external DNS in some apiservers (e.g., <b>kube-apiserver</b> cannot resolve in-cluster DNS as that would be a layering violation). <b>host</b> may also be an IP address.</p> <p>Please note that using <b>localhost</b> or <b>127.0.0.1</b> as a <b>host</b> is risky unless you take great care to run this webhook on all hosts which run an apiserver which might need to make calls to this webhook. Such installs are likely to be non-portable, i.e., not easy to turn up in a new cluster.</p> <p>The scheme must be "https"; the URL must begin with "https://".</p> <p>A path is optional, and if present may be any string permissible in a URL. You may use the path to pass an arbitrary string to the webhook, for example, a cluster identifier.</p> <p>Attempting to use a user or basic auth e.g. "user:password@" is not allowed. Fragments ("#...") and query parameters ("?...") are not allowed, either.</p>

### 5.2.1.5. .spec.conversion.webhook.clientConfig.service

#### Description

ServiceReference holds a reference to Service.legacy.k8s.io

#### Type

**object**

#### Required

- **namespace**

- `namespace`

- `name`

Property	Type	Description
<code>name</code>	<code>string</code>	<code>name</code> is the name of the service. Required
<code>namespace</code>	<code>string</code>	<code>namespace</code> is the namespace of the service. Required
<code>path</code>	<code>string</code>	<code>path</code> is an optional URL path at which the webhook will be contacted.
<code>port</code>	<code>integer</code>	<code>port</code> is an optional service port at which the webhook will be contacted. <b>port</b> should be a valid port number (1-65535, inclusive). Defaults to 443 for backward compatibility.

#### 5.2.1.6. `.spec.names`

##### Description

`CustomResourceDefinitionNames` indicates the names to serve this `CustomResourceDefinition`

##### Type

`object`

##### Required

- `plural`
- `kind`

Property	Type	Description
<code>categories</code>	<code>array (string)</code>	<code>categories</code> is a list of grouped resources this custom resource belongs to (e.g. 'all'). This is published in API discovery documents, and used by clients to support invocations like <b>kubectl get all</b> .

Property	Type	Description
<b>kind</b>	<b>string</b>	kind is the serialized kind of the resource. It is normally CamelCase and singular. Custom resource instances will use this value as the <b>kind</b> attribute in API calls.
<b>listKind</b>	<b>string</b>	listKind is the serialized kind of the list for this resource. Defaults to "`kind` List".
<b>plural</b>	<b>string</b>	plural is the plural name of the resource to serve. The custom resources are served under <b>/apis/&lt;group&gt;/&lt;version&gt;/.../&lt;plural&gt;</b> . Must match the name of the CustomResourceDefinition (in the form <b>&lt;names.plural&gt;</b> . <b>&lt;group&gt;</b> ). Must be all lowercase.
<b>shortNames</b>	<b>array (string)</b>	shortNames are short names for the resource, exposed in API discovery documents, and used by clients to support invocations like <b>kubectl get &lt;shortname&gt;</b> . It must be all lowercase.
<b>singular</b>	<b>string</b>	singular is the singular name of the resource. It must be all lowercase. Defaults to lowercased <b>kind</b> .

### 5.2.1.7. .spec.versions

#### Description

versions is the list of all API versions of the defined custom resource. Version names are used to compute the order in which served versions are listed in API discovery. If the version string is "kube-like", it will sort above non "kube-like" version strings, which are ordered lexicographically. "Kube-like" versions start with a "v", then are followed by a number (the major version), then optionally the string "alpha" or "beta" and another number (the minor version). These are sorted first by GA > beta > alpha (where GA is a version with no suffix such as beta or alpha), and then by comparing major version, then minor version. An example sorted list of versions: v10, v2, v1, v11beta2, v10beta3, v3beta1, v12alpha1, v11alpha2, foo1, foo10.

#### Type

**array**

### 5.2.1.8. .spec.versions[]

**Description**

CustomResourceDefinitionVersion describes a version for CRD.

**Type**

**object**

**Required**

- **name**
- **served**
- **storage**

Property	Type	Description
<b>additionalPrinterColumns</b>	<b>array</b>	additionalPrinterColumns specifies additional columns returned in Table output. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#receiving-resources-as-tables">https://kubernetes.io/docs/reference/using-api/api-concepts/#receiving-resources-as-tables</a> for details. If no columns are specified, a single column displaying the age of the custom resource is used.
<b>additionalPrinterColumns[]</b>	<b>object</b>	CustomResourceColumnDefinition specifies a column for server side printing.
<b>deprecated</b>	<b>boolean</b>	deprecated indicates this version of the custom resource API is deprecated. When set to true, API requests to this version receive a warning header in the server response. Defaults to false.
<b>deprecationWarning</b>	<b>string</b>	deprecationWarning overrides the default warning returned to API clients. May only be set when <b>deprecated</b> is true. The default warning indicates this version is deprecated and recommends use of the newest served version of equal or greater stability, if one exists.

Property	Type	Description
<b>name</b>	<b>string</b>	name is the version name, e.g. "v1", "v2beta1", etc. The custom resources are served under this version at <b>/apis/&lt;group&gt;/&lt;version&gt;/...</b> if <b>served</b> is true.
<b>schema</b>	<b>object</b>	CustomResourceValidation is a list of validation methods for CustomResources.
<b>served</b>	<b>boolean</b>	served is a flag enabling/disabling this version from being served via REST APIs
<b>storage</b>	<b>boolean</b>	storage indicates this version should be used when persisting custom resources to storage. There must be exactly one version with storage=true.
<b>subresources</b>	<b>object</b>	CustomResourceSubresources defines the status and scale subresources for CustomResources.

### 5.2.1.9. .spec.versions[].additionalPrinterColumns

#### Description

additionalPrinterColumns specifies additional columns returned in Table output. See <https://kubernetes.io/docs/reference/using-api/api-concepts/#receiving-resources-as-tables> for details. If no columns are specified, a single column displaying the age of the custom resource is used.

#### Type

**array**

### 5.2.1.10. .spec.versions[].additionalPrinterColumns[]

#### Description

CustomResourceColumnDefinition specifies a column for server side printing.

#### Type

**object**

#### Required

- **name**
- **type**



- `jsonPath`

Property	Type	Description
<b>description</b>	<b>string</b>	description is a human readable description of this column.
<b>format</b>	<b>string</b>	format is an optional OpenAPI type definition for this column. The 'name' format is applied to the primary identifier column to assist in clients identifying column is the resource name. See <a href="https://github.com/OAI/OpenAPI-Specification/blob/master/versions/2.0.md#data-types">https://github.com/OAI/OpenAPI-Specification/blob/master/versions/2.0.md#data-types</a> for details.
<b>jsonPath</b>	<b>string</b>	jsonPath is a simple JSON path (i.e. with array notation) which is evaluated against each custom resource to produce the value for this column.
<b>name</b>	<b>string</b>	name is a human readable name for the column.
<b>priority</b>	<b>integer</b>	priority is an integer defining the relative importance of this column compared to others. Lower numbers are considered higher priority. Columns that may be omitted in limited space scenarios should be given a priority greater than 0.
<b>type</b>	<b>string</b>	type is an OpenAPI type definition for this column. See <a href="https://github.com/OAI/OpenAPI-Specification/blob/master/versions/2.0.md#data-types">https://github.com/OAI/OpenAPI-Specification/blob/master/versions/2.0.md#data-types</a> for details.

### 5.2.1.11. `.spec.versions[].schema`

#### Description

CustomResourceValidation is a list of validation methods for CustomResources.

#### Type

**object**

Property	Type	Description
<b>openAPIV3Schema</b>	<a href="#">JSONSchemaProps</a>	openAPIV3Schema is the OpenAPI v3 schema to use for validation and pruning.

### 5.2.1.12. .spec.versions[].subresources

#### Description

CustomResourceSubresources defines the status and scale subresources for CustomResources.

#### Type

**object**

Property	Type	Description
<b>scale</b>	<b>object</b>	CustomResourceSubresourceScale defines how to serve the scale subresource for CustomResources.
<b>status</b>	<b>object</b>	CustomResourceSubresourceStatus defines how to serve the status subresource for CustomResources. Status is represented by the <b>.status</b> JSON path inside of a CustomResource. When set, * exposes a /status subresource for the custom resource * PUT requests to the /status subresource take a custom resource object, and ignore changes to anything except the status stanza * PUT/POST/PATCH requests to the custom resource ignore changes to the status stanza

### 5.2.1.13. .spec.versions[].subresources.scale

#### Description

CustomResourceSubresourceScale defines how to serve the scale subresource for CustomResources.

#### Type

**object**

#### Required

- **specReplicasPath**
- **statusReplicasPath**

Property	Type	Description
<b>labelSelectorPath</b>	<b>string</b>	labelSelectorPath defines the JSON path inside of a custom resource that corresponds to Scale <b>status.selector</b> . Only JSON paths without the array notation are allowed. Must be a JSON Path under <b>.status</b> or <b>.spec</b> . Must be set to work with HorizontalPodAutoscaler. The field pointed by this JSON path must be a string field (not a complex selector struct) which contains a serialized label selector in string form. More info: <a href="https://kubernetes.io/docs/tasks/access-kubernetes-api/custom-resources/custom-resource-definitions#scale-subresource">https://kubernetes.io/docs/tasks/access-kubernetes-api/custom-resources/custom-resource-definitions#scale-subresource</a> If there is no value under the given path in the custom resource, the <b>status.selector</b> value in the <b>/scale</b> subresource will default to the empty string.
<b>specReplicasPath</b>	<b>string</b>	specReplicasPath defines the JSON path inside of a custom resource that corresponds to Scale <b>spec.replicas</b> . Only JSON paths without the array notation are allowed. Must be a JSON Path under <b>.spec</b> . If there is no value under the given path in the custom resource, the <b>/scale</b> subresource will return an error on GET.
<b>statusReplicasPath</b>	<b>string</b>	statusReplicasPath defines the JSON path inside of a custom resource that corresponds to Scale <b>status.replicas</b> . Only JSON paths without the array notation are allowed. Must be a JSON Path under <b>.status</b> . If there is no value under the given path in the custom resource, the <b>status.replicas</b> value in the <b>/scale</b> subresource will default to 0.

#### 5.2.1.14. .spec.versions[].subresources.status

## Description

CustomResourceSubresourceStatus defines how to serve the status subresource for CustomResources. Status is represented by the **.status** JSON path inside of a CustomResource. When set, \* exposes a /status subresource for the custom resource \* PUT requests to the /status subresource take a custom resource object, and ignore changes to anything except the status stanza \* PUT/POST/PATCH requests to the custom resource ignore changes to the status stanza

## Type

**object**

### 5.2.1.15. .status

## Description

CustomResourceDefinitionStatus indicates the state of the CustomResourceDefinition

## Type

**object**

Property	Type	Description
<b>acceptedNames</b>	<b>object</b>	CustomResourceDefinitionNames indicates the names to serve this CustomResourceDefinition
<b>conditions</b>	<b>array</b>	conditions indicate state for particular aspects of a CustomResourceDefinition
<b>conditions[]</b>	<b>object</b>	CustomResourceDefinitionCondition contains details for the current condition of this pod.
<b>storedVersions</b>	<b>array (string)</b>	storedVersions lists all versions of CustomResources that were ever persisted. Tracking these versions allows a migration path for stored versions in etcd. The field is mutable so a migration controller can finish a migration to another version (ensuring no old objects are left in storage), and then remove the rest of the versions from this list. Versions may not be removed from <b>spec.versions</b> while they exist in this list.

### 5.2.1.16. .status.acceptedNames

## Description

CustomResourceDefinitionNames indicates the names to serve this CustomResourceDefinition

## Type

**object**

## Required

- plural
- kind

Property	Type	Description
<b>categories</b>	<b>array (string)</b>	categories is a list of grouped resources this custom resource belongs to (e.g. 'all'). This is published in API discovery documents, and used by clients to support invocations like <b>kubectl get all</b> .
<b>kind</b>	<b>string</b>	kind is the serialized kind of the resource. It is normally CamelCase and singular. Custom resource instances will use this value as the <b>kind</b> attribute in API calls.
<b>listKind</b>	<b>string</b>	listKind is the serialized kind of the list for this resource. Defaults to "`kind` List".
<b>plural</b>	<b>string</b>	plural is the plural name of the resource to serve. The custom resources are served under <b>/apis/&lt;group&gt;/&lt;version&gt;/.../&lt;plural&gt;</b> . Must match the name of the CustomResourceDefinition (in the form <b>&lt;names.plural&gt;.&lt;group&gt;</b> ). Must be all lowercase.
<b>shortNames</b>	<b>array (string)</b>	shortNames are short names for the resource, exposed in API discovery documents, and used by clients to support invocations like <b>kubectl get &lt;shortname&gt;</b> . It must be all lowercase.
<b>singular</b>	<b>string</b>	singular is the singular name of the resource. It must be all lowercase. Defaults to lowercased <b>kind</b> .

## 5.2.1.17. .status.conditions

## Description

conditions indicate state for particular aspects of a CustomResourceDefinition

Type

**array**

### 5.2.1.18. .status.conditions[]

Description

CustomResourceDefinitionCondition contains details for the current condition of this pod.

Type

**object**

Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	lastTransitionTime last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	message is a human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	reason is a unique, one-word, CamelCase reason for the condition's last transition.
<b>status</b>	<b>string</b>	status is the status of the condition. Can be True, False, Unknown.
<b>type</b>	<b>string</b>	type is the type of the condition. Types include Established, NamesAccepted and Terminating.

### 5.2.2. API endpoints

The following API endpoints are available:

- **/apis/apiextensions.k8s.io/v1/customresourcedefinitions**
  - **DELETE**: delete collection of CustomResourceDefinition
  - **GET**: list or watch objects of kind CustomResourceDefinition
  - **POST**: create a CustomResourceDefinition

- **/apis/apiextensions.k8s.io/v1/watch/customresourcedefinitions**
  - **GET**: watch individual changes to a list of CustomResourceDefinition. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apiextensions.k8s.io/v1/customresourcedefinitions/{name}**
  - **DELETE**: delete a CustomResourceDefinition
  - **GET**: read the specified CustomResourceDefinition
  - **PATCH**: partially update the specified CustomResourceDefinition
  - **PUT**: replace the specified CustomResourceDefinition
- **/apis/apiextensions.k8s.io/v1/watch/customresourcedefinitions/{name}**
  - **GET**: watch changes to an object of kind CustomResourceDefinition. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/apiextensions.k8s.io/v1/customresourcedefinitions/{name}/status**
  - **GET**: read status of the specified CustomResourceDefinition
  - **PATCH**: partially update status of the specified CustomResourceDefinition
  - **PUT**: replace status of the specified CustomResourceDefinition

### 5.2.2.1. /apis/apiextensions.k8s.io/v1/customresourcedefinitions

Table 5.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of CustomResourceDefinition

Table 5.2. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>

Parameter	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 5.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 5.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind CustomResourceDefinition

Table 5.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 5.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CustomResourceDefinitionList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a CustomResourceDefinition

**Table 5.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 5.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CustomResourceDefinition</b> schema	

**Table 5.9. HTTP responses**

HTTP code	Response body
200 - OK	<b>CustomResourceDefinition</b> schema

HTTP code	Response body
201 - Created	<b>CustomResourceDefinition</b> schema
202 - Accepted	<b>CustomResourceDefinition</b> schema
401 - Unauthorized	Empty

### 5.2.2.2. /apis/apiextensions.k8s.io/v1/watch/customresourcedefinitions

Table 5.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.



Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CustomResourceDefinition. deprecated: use the 'watch' parameter with a list operation instead.

Table 5.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 5.2.2.3. /apis/apiextensions.k8s.io/v1/customresourcedefinitions/{name}

Table 5.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CustomResourceDefinition

Table 5.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a CustomResourceDefinition

Table 5.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 5.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 5.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified CustomResourceDefinition

Table 5.17. HTTP responses

HTTP code	Response body
200 - OK	<b>CustomResourceDefinition</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified CustomResourceDefinition

**Table 5.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 5.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 5.20. HTTP responses

HTTP code	Response body
200 - OK	<b>CustomResourceDefinition</b> schema
201 - Created	<b>CustomResourceDefinition</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified CustomResourceDefinition

Table 5.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 5.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CustomResourceDefinition</b> schema	

Table 5.23. HTTP responses

HTTP code	Response body
200 - OK	<b>CustomResourceDefinition</b> schema
201 - Created	<b>CustomResourceDefinition</b> schema
401 - Unauthorized	Empty

#### 5.2.2.4. /apis/apiextensions.k8s.io/v1/watch/customresourcedefinitions/{name}

Table 5.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CustomResourceDefinition

Table 5.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind CustomResourceDefinition. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 5.26. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 5.2.2.5. /apis/apiextensions.k8s.io/v1/customresourcedefinitions/{name}/status

Table 5.27. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CustomResourceDefinition

Table 5.28. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **GET**

#### Description

read status of the specified CustomResourceDefinition

Table 5.29. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">CustomResourceDefinition</a> schema
401 - Unauthorized	Empty

#### HTTP method

#### **PATCH**

#### Description

partially update status of the specified CustomResourceDefinition

Table 5.30. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 5.31. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 5.32. HTTP responses

HTTP code	Response body
200 - OK	<b>CustomResourceDefinition</b> schema
201 - Created	<b>CustomResourceDefinition</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified CustomResourceDefinition

**Table 5.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 5.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<a href="#">CustomResourceDefinition</a> schema	

Table 5.35. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CustomResourceDefinition</a> schema
201 - Created	<a href="#">CustomResourceDefinition</a> schema
401 - Unauthorized	Empty

## CHAPTER 6. REGISTRATION APIS

### 6.1. API REGISTRATION APIS

#### 6.1.1. APIService [apiregistration.k8s.io/v1]

##### Description

APIService represents a server for a particular GroupVersion. Name must be "version.group".

##### Type

**object**

### 6.2. APISERVICE [APIREGISTRATION.K8S.IO/V1]

##### Description

APIService represents a server for a particular GroupVersion. Name must be "version.group".

##### Type

**object**

#### 6.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	APIServiceSpec contains information for locating and communicating with a server. Only https is supported, though you are able to disable certificate verification.
<b>status</b>	<b>object</b>	APIServiceStatus contains derived information about an API server

### 6.2.1.1. .spec

#### Description

APIServiceSpec contains information for locating and communicating with a server. Only https is supported, though you are able to disable certificate verification.

#### Type

**object**

#### Required

- **groupPriorityMinimum**
- **versionPriority**

Property	Type	Description
<b>caBundle</b>	<b>string</b>	CABundle is a PEM encoded CA bundle which will be used to validate an API server's serving certificate. If unspecified, system trust roots on the apiserver are used.
<b>group</b>	<b>string</b>	Group is the API group name this server hosts



Property	Type	Description
<b>groupPriorityMinimum</b>	<b>integer</b>	<p>GroupPriorityMinimum is the priority this group should have at least. Higher priority means that the group is preferred by clients over lower priority ones. Note that other versions of this group might specify even higher GroupPriorityMinimum values such that the whole group gets a higher priority. The primary sort is based on GroupPriorityMinimum, ordered highest number to lowest (20 before 10). The secondary sort is based on the alphabetical comparison of the name of the object. (v1.bar before v1.foo) We'd recommend something like: *.k8s.io (except extensions) at 18000 and PaaSes (OpenShift, Deis) are recommended to be in the 2000s</p>
<b>insecureSkipTLSVerify</b>	<b>boolean</b>	<p>InsecureSkipTLSVerify disables TLS certificate verification when communicating with this server. This is strongly discouraged. You should use the CABundle instead.</p>
<b>service</b>	<b>object</b>	<p>ServiceReference holds a reference to Service.legacy.k8s.io</p>
<b>version</b>	<b>string</b>	<p>Version is the API version this server hosts. For example, "v1"</p>

Property	Type	Description
<b>versionPriority</b>	<b>integer</b>	VersionPriority controls the ordering of this API version inside of its group. Must be greater than zero. The primary sort is based on VersionPriority, ordered highest to lowest (20 before 10). Since it's inside of a group, the number can be small, probably in the 10s. In case of equal version priorities, the version string will be used to compute the order inside a group. If the version string is "kube-like", it will sort above non "kube-like" version strings, which are ordered lexicographically. "Kube-like" versions start with a "v", then are followed by a number (the major version), then optionally the string "alpha" or "beta" and another number (the minor version). These are sorted first by GA > beta > alpha (where GA is a version with no suffix such as beta or alpha), and then by comparing major version, then minor version. An example sorted list of versions: v10, v2, v1, v11beta2, v10beta3, v3beta1, v12alpha1, v11alpha2, foo1, foo10.

### 6.2.1.2. .spec.service

#### Description

ServiceReference holds a reference to Service.legacy.k8s.io

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name is the name of the service
<b>namespace</b>	<b>string</b>	Namespace is the namespace of the service

Property	Type	Description
<b>port</b>	<b>integer</b>	If specified, the port on the service that hosting webhook. Default to 443 for backward compatibility. <b>port</b> should be a valid port number (1-65535, inclusive).

### 6.2.1.3. .status

#### Description

APIServiceStatus contains derived information about an API server

#### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array</b>	Current service state of apiService.
<b>conditions[]</b>	<b>object</b>	APIServiceCondition describes the state of an APIService at a particular point

### 6.2.1.4. .status.conditions

#### Description

Current service state of apiService.

#### Type

**array**

### 6.2.1.5. .status.conditions[]

#### Description

APIServiceCondition describes the state of an APIService at a particular point

#### Type

**object**

#### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	Human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	Unique, one-word, CamelCase reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status is the status of the condition. Can be True, False, Unknown.
<b>type</b>	<b>string</b>	Type is the type of the condition.

## 6.2.2. API endpoints

The following API endpoints are available:

- **/apis/apiregistration.k8s.io/v1/apiservices**
  - **DELETE**: delete collection of APIService
  - **GET**: list or watch objects of kind APIService
  - **POST**: create an APIService
- **/apis/apiregistration.k8s.io/v1/watch/apiservices**
  - **GET**: watch individual changes to a list of APIService. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apiregistration.k8s.io/v1/apiservices/{name}**
  - **DELETE**: delete an APIService
  - **GET**: read the specified APIService
  - **PATCH**: partially update the specified APIService
  - **PUT**: replace the specified APIService
- **/apis/apiregistration.k8s.io/v1/watch/apiservices/{name}**
  - **GET**: watch changes to an object of kind APIService. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/apiregistration.k8s.io/v1/apiservices/{name}/status**

- **GET**: read status of the specified APIService
- **PATCH**: partially update status of the specified APIService
- **PUT**: replace status of the specified APIService

### 6.2.2.1. /apis/apiregistration.k8s.io/v1/apiservices

Table 6.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of APIService

Table 6.2. Query parameters

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 6.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 6.4. HTTP responses



HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind APIService

**Table 6.5. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 6.6. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">APIServiceList</a> schema
401 - Unauthorized	Empty

## HTTP method

**POST****Description**

create an APIService

**Table 6.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 6.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>APIService</b> schema	

**Table 6.9. HTTP responses**

HTTP code	Response body
200 - OK	<b>APIService</b> schema

HTTP code	Response body
201 - Created	<b>APIService</b> schema
202 - Accepted	<b>APIService</b> schema
401 - Unauthorized	Empty

### 6.2.2.2. /apis/apiregistration.k8s.io/v1/watch/apiservices

Table 6.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of APIService. deprecated: use the 'watch' parameter with a list operation instead.

Table 6.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 6.2.2.3. /apis/apiregistration.k8s.io/v1/apiservices/{name}

Table 6.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the APIService

Table 6.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete an APIService

Table 6.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 6.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 6.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified APIService

Table 6.17. HTTP responses

HTTP code	Response body
200 - OK	<b>APIService</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified APIService

**Table 6.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 6.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 6.20. HTTP responses

HTTP code	Reponse body
200 - OK	<b>APIService</b> schema
201 - Created	<b>APIService</b> schema
401 - Unauthorized	Empty

## HTTP method

### PUT

#### Description

replace the specified APIService

Table 6.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 6.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>APIService</b> schema	

Table 6.23. HTTP responses

HTTP code	Response body
200 - OK	<b>APIService</b> schema
201 - Created	<b>APIService</b> schema
401 - Unauthorized	Empty

#### 6.2.2.4. /apis/apiregistration.k8s.io/v1/watch/apiservices/{name}

Table 6.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the APIService

Table 6.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind APIService. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 6.26. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 6.2.2.5. /apis/apiregistration.k8s.io/v1/apiservices/{name}/status

Table 6.27. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the APIService

Table 6.28. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### GET

#### Description

read status of the specified APIService

Table 6.29. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">APIService</a> schema
401 - Unauthorized	Empty

#### HTTP method

#### PATCH

#### Description

partially update status of the specified APIService

Table 6.30. Query parameters

Parameter	Type	Description
-----------	------	-------------



Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 6.31. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 6.32. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">APIService</a> schema
201 - Created	<a href="#">APIService</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified [APIService](#)

**Table 6.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 6.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>APIService</b> schema	

Table 6.35. HTTP responses

HTTP code	Response body
200 - OK	<b>APIService</b> schema
201 - Created	<b>APIService</b> schema
401 - Unauthorized	Empty

## CHAPTER 7. APPS APIS

### 7.1. APPS APIS

#### 7.1.1. ControllerRevision [apps/v1]

##### Description

ControllerRevision implements an immutable snapshot of state data. Clients are responsible for serializing and deserializing the objects that contain their internal state. Once a ControllerRevision has been successfully created, it can not be updated. The API Server will fail validation of all requests that attempt to mutate the Data field. ControllerRevisions may, however, be deleted. Note that, due to its use by both the DaemonSet and StatefulSet controllers for update and rollback, this object is beta. However, it may be subject to name and representation changes in future releases, and clients should not depend on its stability. It is primarily for internal use by controllers.

##### Type

**object**

#### 7.1.2. DaemonSet [apps/v1]

##### Description

DaemonSet represents the configuration of a daemon set.

##### Type

**object**

#### 7.1.3. Deployment [apps/v1]

##### Description

Deployment enables declarative updates for Pods and ReplicaSets.

##### Type

**object**

#### 7.1.4. ReplicaSet [apps/v1]

##### Description

ReplicaSet ensures that a specified number of pod replicas are running at any given time.

##### Type

**object**

#### 7.1.5. StatefulSet [apps/v1]

##### Description

StatefulSet represents a set of pods with consistent identities. Identities are defined as: - Network: A single stable DNS and hostname. - Storage: As many VolumeClaims as requested.

The StatefulSet guarantees that a given network identity will always map to the same storage identity.

##### Type

**object**

## 7.2. CONTROLLERREVISION [APPS/V1]

### Description

ControllerRevision implements an immutable snapshot of state data. Clients are responsible for serializing and deserializing the objects that contain their internal state. Once a ControllerRevision has been successfully created, it can not be updated. The API Server will fail validation of all requests that attempt to mutate the Data field. ControllerRevisions may, however, be deleted. Note that, due to its use by both the DaemonSet and StatefulSet controllers for update and rollback, this object is beta. However, it may be subject to name and representation changes in future releases, and clients should not depend on its stability. It is primarily for internal use by controllers.

### Type

**object**

### Required

- **revision**

### 7.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>data</b>	<b>RawExtension</b>	Data is the serialized representation of the state.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>revision</b>	<b>integer</b>	Revision indicates the revision of the state represented by Data.

## 7.2.2. API endpoints

The following API endpoints are available:

- **/apis/apps/v1/controllerrevisions**
  - **GET**: list or watch objects of kind ControllerRevision
- **/apis/apps/v1/watch/controllerrevisions**
  - **GET**: watch individual changes to a list of ControllerRevision. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/controllerrevisions**
  - **DELETE**: delete collection of ControllerRevision
  - **GET**: list or watch objects of kind ControllerRevision
  - **POST**: create a ControllerRevision
- **/apis/apps/v1/watch/namespaces/{namespace}/controllerrevisions**
  - **GET**: watch individual changes to a list of ControllerRevision. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/controllerrevisions/{name}**
  - **DELETE**: delete a ControllerRevision
  - **GET**: read the specified ControllerRevision
  - **PATCH**: partially update the specified ControllerRevision
  - **PUT**: replace the specified ControllerRevision
- **/apis/apps/v1/watch/namespaces/{namespace}/controllerrevisions/{name}**
  - **GET**: watch changes to an object of kind ControllerRevision. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 7.2.2.1. /apis/apps/v1/controllerrevisions

Table 7.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind ControllerRevision

Table 7.2. HTTP responses

HTTP code	Response body
200 - OK	<b>ControllerRevisionList</b> schema
401 - Unauthorized	Empty

### 7.2.2.2. /apis/apps/v1/watch/controllerrevisions

Table 7.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ControllerRevision. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.4. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.2.2.3. /apis/apps/v1/namespaces/{namespace}/controllerrevisions

Table 7.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ControllerRevision

Table 7.7. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>



Table 7.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.9. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ControllerRevision

Table 7.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 7.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ControllerRevisionList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a ControllerRevision

**Table 7.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ControllerRevision</b> schema	

**Table 7.14. HTTP responses**

HTTP code	Response body
200 - OK	<b>ControllerRevision</b> schema

HTTP code	Response body
201 - Created	<a href="#">ControllerRevision</a> schema
202 - Accepted	<a href="#">ControllerRevision</a> schema
401 - Unauthorized	Empty

#### 7.2.2.4. /apis/apps/v1/watch/namespaces/{namespace}/controllerrevisions

Table 7.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ControllerRevision. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.2.2.5. /apis/apps/v1/namespaces/{namespace}/controllerrevisions/{name}

Table 7.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ControllerRevision
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a ControllerRevision

Table 7.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 7.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ControllerRevision

Table 7.23. HTTP responses

HTTP code	Response body
200 - OK	<b>ControllerRevision</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ControllerRevision

**Table 7.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.26. HTTP responses

HTTP code	Reponse body
200 - OK	<b>ControllerRevision</b> schema
201 - Created	<b>ControllerRevision</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ControllerRevision

Table 7.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 7.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ControllerRevision</b> schema	

Table 7.29. HTTP responses

HTTP code	Response body
200 - OK	<b>ControllerRevision</b> schema
201 - Created	<b>ControllerRevision</b> schema
401 - Unauthorized	Empty

#### 7.2.2.6. /apis/apps/v1/watch/namespaces/{namespace}/controllerrevisions/{name}

Table 7.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ControllerRevision
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 7.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ControllerRevision. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 7.32. HTTP responses**

HTTP code	Response body
200 - OK	<b>WatchEvent</b> schema
401 - Unauthorized	Empty

## 7.3. DAEMONSET [APPS/V1]

### Description

DaemonSet represents the configuration of a daemon set.

### Type

**object**

### 7.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	DaemonSetSpec is the specification of a daemon set.

Property	Type	Description
<b>status</b>	<b>object</b>	DaemonSetStatus represents the current status of a daemon set.

### 7.3.1.1. .spec

#### Description

DaemonSetSpec is the specification of a daemon set.

#### Type

**object**

#### Required

- **selector**
- **template**

Property	Type	Description
<b>minReadySeconds</b>	<b>integer</b>	The minimum number of seconds for which a newly created DaemonSet pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready).
<b>revisionHistoryLimit</b>	<b>integer</b>	The number of old history to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
<b>selector</b>	<b>LabelSelector</b>	A label query over pods that are managed by the daemon set. Must match in order to be controlled. It must match the pod template's labels. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors</a>

Property	Type	Description
<b>template</b>	<b>PodTemplateSpec</b>	An object that describes the pod that will be created. The DaemonSet will create exactly one copy of this pod on every node that matches the template's node selector (or on every node if no node selector is specified). The only allowed template.spec.restartPolicy value is "Always". More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#pod-template">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#pod-template</a>
<b>updateStrategy</b>	<b>object</b>	DaemonSetUpdateStrategy is a struct used to control the update strategy for a DaemonSet.

### 7.3.1.2. .spec.updateStrategy

#### Description

DaemonSetUpdateStrategy is a struct used to control the update strategy for a DaemonSet.

#### Type

**object**

Property	Type	Description
<b>rollingUpdate</b>	<b>object</b>	Spec to control the desired behavior of daemon set rolling update.
<b>type</b>	<b>string</b>	Type of daemon set update. Can be "RollingUpdate" or "OnDelete". Default is RollingUpdate.  Possible enum values: - <b>"OnDelete"</b> Replace the old daemons only when it's killed - <b>"RollingUpdate"</b> Replace the old daemons by new ones using rolling update i.e replace them on each node one after the other.

### 7.3.1.3. .spec.updateStrategy.rollingUpdate

#### Description

Spec to control the desired behavior of daemon set rolling update.

Type  
object

Property	Type	Description
<b>maxSurge</b>	<b>IntOrString</b>	<p>The maximum number of nodes with an existing available DaemonSet pod that can have an updated DaemonSet pod during during an update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up to a minimum of 1. Default value is 0. Example: when this is set to 30%, at most 30% of the total number of nodes that should be running the daemon pod (i.e. <code>status.desiredNumberScheduled</code>) can have their a new pod created before the old pod is marked as deleted. The update starts by launching new pods on 30% of nodes. Once an updated pod is available (Ready for at least <code>minReadySeconds</code>) the old DaemonSet pod on that node is marked deleted. If the old pod becomes unavailable for any reason (Ready transitions to false, is evicted, or is drained) an updated pod is immediately created on that node without considering surge limits. Allowing surge implies the possibility that the resources consumed by the daemonset on any given node can double if the readiness check fails, and so resource intensive daemonsets should take into account that they may cause evictions during disruption.</p>

Property	Type	Description
<b>maxUnavailable</b>	<b>IntOrString</b>	The maximum number of DaemonSet pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of total number of DaemonSet pods at the start of the update (ex: 10%). Absolute number is calculated from percentage by rounding up. This cannot be 0 if MaxSurge is 0 Default value is 1. Example: when this is set to 30%, at most 30% of the total number of nodes that should be running the daemon pod (i.e. status.desiredNumberScheduled) can have their pods stopped for an update at any given time. The update starts by stopping at most 30% of those DaemonSet pods and then brings up new DaemonSet pods in their place. Once the new pods are available, it then proceeds onto other DaemonSet pods, thus ensuring that at least 70% of original number of DaemonSet pods are available at all times during the update.

#### 7.3.1.4. .status

##### Description

DaemonSetStatus represents the current status of a daemon set.

##### Type

**object**

##### Required

- **currentNumberScheduled**
- **numberMisscheduled**
- **desiredNumberScheduled**
- **numberReady**

Property	Type	Description
<b>collisionCount</b>	<b>integer</b>	Count of hash collisions for the DaemonSet. The DaemonSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
<b>conditions</b>	<b>array</b>	Represents the latest available observations of a DaemonSet's current state.
<b>conditions[]</b>	<b>object</b>	DaemonSetCondition describes the state of a DaemonSet at a certain point.
<b>currentNumberScheduled</b>	<b>integer</b>	The number of nodes that are running at least 1 daemon pod and are supposed to run the daemon pod. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/">https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/</a>
<b>desiredNumberScheduled</b>	<b>integer</b>	The total number of nodes that should be running the daemon pod (including nodes correctly running the daemon pod). More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/">https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/</a>
<b>numberAvailable</b>	<b>integer</b>	The number of nodes that should be running the daemon pod and have one or more of the daemon pod running and available (ready for at least spec.minReadySeconds)
<b>numberMisscheduled</b>	<b>integer</b>	The number of nodes that are running the daemon pod, but are not supposed to run the daemon pod. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/">https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/</a>

Property	Type	Description
<b>numberReady</b>	<b>integer</b>	numberReady is the number of nodes that should be running the daemon pod and have one or more of the daemon pod running with a Ready Condition.
<b>numberUnavailable</b>	<b>integer</b>	The number of nodes that should be running the daemon pod and have none of the daemon pod running and available (ready for at least spec.minReadySeconds)
<b>observedGeneration</b>	<b>integer</b>	The most recent generation observed by the daemon set controller.
<b>updatedNumberScheduled</b>	<b>integer</b>	The total number of nodes that are running updated daemon pod

### 7.3.1.5. .status.conditions

#### Description

Represents the latest available observations of a DaemonSet's current state.

#### Type

**array**

### 7.3.1.6. .status.conditions[]

#### Description

DaemonSetCondition describes the state of a DaemonSet at a certain point.

#### Type

**object**

#### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transitioned from one status to another.



Property	Type	Description
<b>message</b>	<b>string</b>	A human readable message indicating details about the transition.
<b>reason</b>	<b>string</b>	The reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of DaemonSet condition.

### 7.3.2. API endpoints

The following API endpoints are available:

- **/apis/apps/v1/daemonsets**
  - **GET**: list or watch objects of kind DaemonSet
- **/apis/apps/v1/watch/daemonsets**
  - **GET**: watch individual changes to a list of DaemonSet. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/daemonsets**
  - **DELETE**: delete collection of DaemonSet
  - **GET**: list or watch objects of kind DaemonSet
  - **POST**: create a DaemonSet
- **/apis/apps/v1/watch/namespaces/{namespace}/daemonsets**
  - **GET**: watch individual changes to a list of DaemonSet. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/daemonsets/{name}**
  - **DELETE**: delete a DaemonSet
  - **GET**: read the specified DaemonSet
  - **PATCH**: partially update the specified DaemonSet
  - **PUT**: replace the specified DaemonSet
- **/apis/apps/v1/watch/namespaces/{namespace}/daemonsets/{name}**
  - **GET**: watch changes to an object of kind DaemonSet. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector'

parameter.

- **/apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status**
  - **GET**: read status of the specified DaemonSet
  - **PATCH**: partially update status of the specified DaemonSet
  - **PUT**: replace status of the specified DaemonSet

### 7.3.2.1. /apis/apps/v1/daemonsets

Table 7.33. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind DaemonSet

Table 7.34. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">DaemonSetList</a> schema
401 - Unauthorized	Empty

### 7.3.2.2. /apis/apps/v1/watch/daemonsets

Table 7.35. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of DaemonSet. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.36. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.3.2.3. /apis/apps/v1/namespaces/{namespace}/daemonsets

Table 7.37. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.38. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of DaemonSet

Table 7.39. Query parameters

Parameter	Type	Description
-----------	------	-------------



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 7.40. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.41. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind DaemonSet

Table 7.42. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 7.43. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">DaemonSetList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a DaemonSet

**Table 7.44. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.45. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>DaemonSet</b> schema	

**Table 7.46. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>DaemonSet</b> schema



HTTP code	Response body
201 - Created	<b>DaemonSet</b> schema
202 - Accepted	<b>DaemonSet</b> schema
401 - Unauthorized	Empty

### 7.3.2.4. /apis/apps/v1/watch/namespaces/{namespace}/daemonsets

Table 7.47. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.48. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of DaemonSet. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.49. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.3.2.5. /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}

Table 7.50. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the DaemonSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.51. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a DaemonSet

Table 7.52. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 7.53. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.54. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified DaemonSet

Table 7.55. HTTP responses

HTTP code	Response body
200 - OK	<b>DaemonSet</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified DaemonSet

**Table 7.56. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.57. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.58. HTTP responses

HTTP code	Response body
200 - OK	<b>DaemonSet</b> schema
201 - Created	<b>DaemonSet</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified DaemonSet

Table 7.59. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 7.60. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DaemonSet</b> schema	

Table 7.61. HTTP responses

HTTP code	Response body
200 - OK	<b>DaemonSet</b> schema
201 - Created	<b>DaemonSet</b> schema
401 - Unauthorized	Empty

### 7.3.2.6. /apis/apps/v1/watch/namespaces/{namespace}/daemonsets/{name}

Table 7.62. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the DaemonSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.63. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind DaemonSet. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 7.64. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.3.2.7. /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status

Table 7.65. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the DaemonSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.66. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified DaemonSet

Table 7.67. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">DaemonSet</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified DaemonSet

Table 7.68. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.69. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.70. HTTP responses

HTTP code	Response body
200 - OK	<b>DaemonSet</b> schema
201 - Created	<b>DaemonSet</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified DaemonSet

**Table 7.71. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.72. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>DaemonSet</b> schema	

Table 7.73. HTTP responses

HTTP code	Response body
200 - OK	<b>DaemonSet</b> schema
201 - Created	<b>DaemonSet</b> schema
401 - Unauthorized	Empty

## 7.4. DEPLOYMENT [APPS/V1]

### Description

Deployment enables declarative updates for Pods and ReplicaSets.

### Type

**object**

### 7.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>



Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	DeploymentSpec is the specification of the desired behavior of the Deployment.
<b>status</b>	<b>object</b>	DeploymentStatus is the most recently observed status of the Deployment.

#### 7.4.1.1. .spec

##### Description

DeploymentSpec is the specification of the desired behavior of the Deployment.

##### Type

**object**

##### Required

- **selector**
- **template**

Property	Type	Description
<b>minReadySeconds</b>	<b>integer</b>	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
<b>paused</b>	<b>boolean</b>	Indicates that the deployment is paused.

Property	Type	Description
<b>progressDeadlineSeconds</b>	<b>integer</b>	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
<b>replicas</b>	<b>integer</b>	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
<b>revisionHistoryLimit</b>	<b>integer</b>	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
<b>selector</b>	<b>LabelSelector</b>	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
<b>strategy</b>	<b>object</b>	DeploymentStrategy describes how to replace existing pods with new ones.
<b>template</b>	<b>PodTemplateSpec</b>	Template describes the pods that will be created. The only allowed template.spec.restartPolicy value is "Always".

### 7.4.1.2. .spec.strategy

#### Description

DeploymentStrategy describes how to replace existing pods with new ones.

#### Type

**object**

Property	Type	Description
<b>rollingUpdate</b>	<b>object</b>	Spec to control the desired behavior of rolling update.
<b>type</b>	<b>string</b>	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.  Possible enum values: - <b>"Recreate"</b> Kill all existing pods before creating new ones. - <b>"RollingUpdate"</b> Replace the old ReplicaSets by new one using rolling update i.e gradually scale down the old ReplicaSets and scale up the new one.

### 7.4.1.3. .spec.strategy.rollingUpdate

#### Description

Spec to control the desired behavior of rolling update.

#### Type

**object**

Property	Type	Description
<b>maxSurge</b>	<b>IntOrString</b>	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

Property	Type	Description
<b>maxUnavailable</b>	<b>IntOrString</b>	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

#### 7.4.1.4. .status

##### Description

DeploymentStatus is the most recently observed status of the Deployment.

##### Type

**object**

Property	Type	Description
<b>availableReplicas</b>	<b>integer</b>	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
<b>collisionCount</b>	<b>integer</b>	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
<b>conditions</b>	<b>array</b>	Represents the latest available observations of a deployment's current state.

Property	Type	Description
<b>conditions[]</b>	<b>object</b>	DeploymentCondition describes the state of a deployment at a certain point.
<b>observedGeneration</b>	<b>integer</b>	The generation observed by the deployment controller.
<b>readyReplicas</b>	<b>integer</b>	readyReplicas is the number of pods targeted by this Deployment with a Ready Condition.
<b>replicas</b>	<b>integer</b>	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
<b>unavailableReplicas</b>	<b>integer</b>	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.
<b>updatedReplicas</b>	<b>integer</b>	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

#### 7.4.1.5. .status.conditions

##### Description

Represents the latest available observations of a deployment's current state.

##### Type

**array**

#### 7.4.1.6. .status.conditions[]

##### Description

DeploymentCondition describes the state of a deployment at a certain point.

##### Type

**object**

##### Required

- **type**

- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transitioned from one status to another.
<b>lastUpdateTime</b>	<b>Time</b>	The last time this condition was updated.
<b>message</b>	<b>string</b>	A human readable message indicating details about the transition.
<b>reason</b>	<b>string</b>	The reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of deployment condition.

## 7.4.2. API endpoints

The following API endpoints are available:

- **/apis/apps/v1/deployments**
  - **GET**: list or watch objects of kind Deployment
- **/apis/apps/v1/watch/deployments**
  - **GET**: watch individual changes to a list of Deployment. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/deployments**
  - **DELETE**: delete collection of Deployment
  - **GET**: list or watch objects of kind Deployment
  - **POST**: create a Deployment
- **/apis/apps/v1/watch/namespaces/{namespace}/deployments**
  - **GET**: watch individual changes to a list of Deployment. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/deployments/{name}**
  - **DELETE**: delete a Deployment

- **GET**: read the specified Deployment
- **PATCH**: partially update the specified Deployment
- **PUT**: replace the specified Deployment
- **/apis/apps/v1/watch/namespaces/{namespace}/deployments/{name}**
  - **GET**: watch changes to an object of kind Deployment. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/apps/v1/namespaces/{namespace}/deployments/{name}/status**
  - **GET**: read status of the specified Deployment
  - **PATCH**: partially update status of the specified Deployment
  - **PUT**: replace status of the specified Deployment

#### 7.4.2.1. /apis/apps/v1/deployments

Table 7.74. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Deployment

**Table 7.75. HTTP responses**

HTTP code	Response body
200 - OK	<b>DeploymentList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 7.4.2.2. /apis/apps/v1/watch/deployments

Table 7.76. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Deployment. deprecated: use the 'watch' parameter with a list operation instead.

**Table 7.77. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.4.2.3. /apis/apps/v1/namespaces/{namespace}/deployments

Table 7.78. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.79. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Deployment

Table 7.80. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 7.81. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.82. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Deployment

Table 7.83. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 7.84. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">DeploymentList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a Deployment

**Table 7.85. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.86. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Deployment</b> schema	

**Table 7.87. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Deployment</b> schema

HTTP code	Reponse body
201 - Created	<b>Deployment</b> schema
202 - Accepted	<b>Deployment</b> schema
401 - Unauthorized	Empty

#### 7.4.2.4. /apis/apps/v1/watch/namespaces/{namespace}/deployments

Table 7.88. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.89. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Deployment. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.90. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 7.4.2.5. /apis/apps/v1/namespaces/{namespace}/deployments/{name}

Table 7.91. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Deployment
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.92. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a Deployment

Table 7.93. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 7.94. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.95. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Deployment

Table 7.96. HTTP responses

HTTP code	Response body
200 - OK	<b>Deployment</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Deployment

**Table 7.97. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.98. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.99. HTTP responses

HTTP code	Response body
200 - OK	<b>Deployment</b> schema
201 - Created	<b>Deployment</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Deployment

Table 7.100. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 7.101. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Deployment</b> schema	

Table 7.102. HTTP responses

HTTP code	Response body
200 - OK	<b>Deployment</b> schema
201 - Created	<b>Deployment</b> schema
401 - Unauthorized	Empty

#### 7.4.2.6. /apis/apps/v1/watch/namespaces/{namespace}/deployments/{name}

Table 7.103. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Deployment
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.104. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Deployment. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 7.105. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 7.4.2.7. /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status

Table 7.106. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Deployment
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.107. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Deployment

Table 7.108. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">Deployment</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Deployment

Table 7.109. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.110. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.111. HTTP responses

HTTP code	Response body
200 - OK	<b>Deployment</b> schema
201 - Created	<b>Deployment</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Deployment

**Table 7.112. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.113. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Deployment</b> schema	

Table 7.114. HTTP responses

HTTP code	Response body
200 - OK	<b>Deployment</b> schema
201 - Created	<b>Deployment</b> schema
401 - Unauthorized	Empty

## 7.5. REPLICASET [APPS/V1]

### Description

ReplicaSet ensures that a specified number of pod replicas are running at any given time.

### Type

**object**

### 7.5.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	If the Labels of a ReplicaSet are empty, they are defaulted to be the same as the Pod(s) that the ReplicaSet manages. Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	ReplicaSetSpec is the specification of a ReplicaSet.
<b>status</b>	<b>object</b>	ReplicaSetStatus represents the current status of a ReplicaSet.

### 7.5.1.1. .spec

#### Description

ReplicaSetSpec is the specification of a ReplicaSet.

#### Type

**object**

#### Required

- **selector**

Property	Type	Description
<b>minReadySeconds</b>	<b>integer</b>	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
<b>replicas</b>	<b>integer</b>	Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller</a>

Property	Type	Description
<b>selector</b>	<a href="#">LabelSelector</a>	Selector is a label query over pods that should match the replica count. Label keys and values that must match in order to be controlled by this replica set. It must match the pod template's labels. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors</a>
<b>template</b>	<a href="#">PodTemplateSpec</a>	Template is the object that describes the pod that will be created if insufficient replicas are detected. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#pod-template">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#pod-template</a>

### 7.5.1.2. .status

#### Description

ReplicaSetStatus represents the current status of a ReplicaSet.

#### Type

**object**

#### Required

- **replicas**

Property	Type	Description
<b>availableReplicas</b>	<b>integer</b>	The number of available replicas (ready for at least <code>minReadySeconds</code> ) for this replica set.
<b>conditions</b>	<b>array</b>	Represents the latest available observations of a replica set's current state.
<b>conditions[]</b>	<b>object</b>	ReplicaSetCondition describes the state of a replica set at a certain point.



Property	Type	Description
<b>fullyLabeledReplicas</b>	<b>integer</b>	The number of pods that have labels matching the labels of the pod template of the replicaset.
<b>observedGeneration</b>	<b>integer</b>	ObservedGeneration reflects the generation of the most recently observed ReplicaSet.
<b>readyReplicas</b>	<b>integer</b>	readyReplicas is the number of pods targeted by this ReplicaSet with a Ready Condition.
<b>replicas</b>	<b>integer</b>	Replicas is the most recently observed number of replicas. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller</a>

### 7.5.1.3. .status.conditions

#### Description

Represents the latest available observations of a replica set's current state.

#### Type

**array**

### 7.5.1.4. .status.conditions[]

#### Description

ReplicaSetCondition describes the state of a replica set at a certain point.

#### Type

**object**

#### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	The last time the condition transitioned from one status to another.

Property	Type	Description
<b>message</b>	<b>string</b>	A human readable message indicating details about the transition.
<b>reason</b>	<b>string</b>	The reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of replica set condition.

## 7.5.2. API endpoints

The following API endpoints are available:

- **/apis/apps/v1/replicasets**
  - **GET**: list or watch objects of kind ReplicaSet
- **/apis/apps/v1/watch/replicasets**
  - **GET**: watch individual changes to a list of ReplicaSet. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/replicasets**
  - **DELETE**: delete collection of ReplicaSet
  - **GET**: list or watch objects of kind ReplicaSet
  - **POST**: create a ReplicaSet
- **/apis/apps/v1/watch/namespaces/{namespace}/replicasets**
  - **GET**: watch individual changes to a list of ReplicaSet. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/replicasets/{name}**
  - **DELETE**: delete a ReplicaSet
  - **GET**: read the specified ReplicaSet
  - **PATCH**: partially update the specified ReplicaSet
  - **PUT**: replace the specified ReplicaSet
- **/apis/apps/v1/watch/namespaces/{namespace}/replicasets/{name}**
  - **GET**: watch changes to an object of kind ReplicaSet. deprecated: use the 'watch' parameter

with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

- **/apis/apps/v1/namespaces/{namespace}/replicasets/{name}/status**
  - **GET**: read status of the specified ReplicaSet
  - **PATCH**: partially update status of the specified ReplicaSet
  - **PUT**: replace status of the specified ReplicaSet

### 7.5.2.1. /apis/apps/v1/replicasets

Table 7.115. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind ReplicaSet

Table 7.116. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ReplicaSetList</a> schema
401 - Unauthorized	Empty

### 7.5.2.2. /apis/apps/v1/watch/replicasets

Table 7.117. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ReplicaSet. deprecated: use the 'watch' parameter with a list operation instead.

**Table 7.118. HTTP responses**



HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.5.2.3. /apis/apps/v1/namespaces/{namespace}/replicasets

Table 7.119. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.120. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ReplicaSet

Table 7.121. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 7.122. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.123. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ReplicaSet

Table 7.124. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 7.125. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ReplicaSetList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**



**Description**

create a ReplicaSet

**Table 7.126. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.127. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ReplicaSet</b> schema	

**Table 7.128. HTTP responses**

HTTP code	Response body
200 - OK	<b>ReplicaSet</b> schema

HTTP code	Response body
201 - Created	<a href="#">ReplicaSet</a> schema
202 - Accepted	<a href="#">ReplicaSet</a> schema
401 - Unauthorized	Empty

#### 7.5.2.4. /apis/apps/v1/watch/namespaces/{namespace}/replicasets

Table 7.129. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.130. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ReplicaSet. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.131. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.5.2.5. /apis/apps/v1/namespaces/{namespace}/replicasets/{name}

Table 7.132. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ReplicaSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.133. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a ReplicaSet

Table 7.134. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 7.135. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.136. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ReplicaSet

Table 7.137. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicaSet</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ReplicaSet

**Table 7.138. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.



Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.139. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.140. HTTP responses

HTTP code	Reponse body
200 - OK	<b>ReplicaSet</b> schema
201 - Created	<b>ReplicaSet</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ReplicaSet

Table 7.141. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 7.142. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ReplicaSet</b> schema	

Table 7.143. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicaSet</b> schema
201 - Created	<b>ReplicaSet</b> schema
401 - Unauthorized	Empty

#### 7.5.2.6. /apis/apps/v1/watch/namespaces/{namespace}/replicasets/{name}

Table 7.144. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ReplicaSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.145. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ReplicaSet. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 7.146. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.5.2.7. /apis/apps/v1/namespaces/{namespace}/replicasets/{name}/status

Table 7.147. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ReplicaSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.148. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified ReplicaSet

Table 7.149. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">ReplicaSet</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified ReplicaSet

Table 7.150. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.151. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.152. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicaSet</b> schema
201 - Created	<b>ReplicaSet</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified ReplicaSet

**Table 7.153. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.154. Body parameters**



Parameter	Type	Description
<b>body</b>	<b>ReplicaSet</b> schema	

Table 7.155. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicaSet</b> schema
201 - Created	<b>ReplicaSet</b> schema
401 - Unauthorized	Empty

## 7.6. STATEFULSET [APPS/V1]

### Description

StatefulSet represents a set of pods with consistent identities. Identities are defined as: - Network: A single stable DNS and hostname. - Storage: As many VolumeClaims as requested.

The StatefulSet guarantees that a given network identity will always map to the same storage identity.

### Type

**object**

### 7.6.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	A StatefulSetSpec is the specification of a StatefulSet.
<b>status</b>	<b>object</b>	StatefulSetStatus represents the current state of a StatefulSet.

### 7.6.1.1. .spec

#### Description

A StatefulSetSpec is the specification of a StatefulSet.

#### Type

**object**

#### Required

- **selector**
- **template**
- **serviceName**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>minReadySeconds</b>	<b>integer</b>	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
<b>ordinals</b>	<b>object</b>	StatefulSetOrdinals describes the policy used for replica ordinal assignment in this StatefulSet.
<b>persistentVolumeClaimRetentionPolicy</b>	<b>object</b>	StatefulSetPersistentVolumeClaimRetentionPolicy describes the policy used for PVCs created from the StatefulSet VolumeClaimTemplates.
<b>podManagementPolicy</b>	<b>string</b>	<p>podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <b>OrderedReady</b>, where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <b>Parallel</b> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.</p> <p>Possible enum values: - <b>"OrderedReady"</b> will create pods in strictly increasing order on scale up and strictly decreasing order on scale down, progressing only when the previous pod is ready or terminated. At most one pod will be changed at any time. - <b>"Parallel"</b> will create and delete pods as soon as the stateful set replica count is changed, and will not wait for pods to be ready or complete termination.</p>

Property	Type	Description
<b>replicas</b>	<b>integer</b>	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
<b>revisionHistoryLimit</b>	<b>integer</b>	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors</a>
<b>serviceName</b>	<b>string</b>	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.

Property	Type	Description
<b>template</b>	<b>PodTemplateSpec</b>	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet. Each pod will be named with the format <statefulsetname>-<podindex>. For example, a pod in a StatefulSet named "web" with index number "3" would be named "web-3". The only allowed template.spec.restartPolicy value is "Always".
<b>updateStrategy</b>	<b>object</b>	StatefulSetUpdateStrategy indicates the strategy that the StatefulSet controller will use to perform updates. It includes any additional parameters necessary to perform the update for the indicated strategy.
<b>volumeClaimTemplates</b>	<b>array (PersistentVolumeClaim)</b>	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

### 7.6.1.2. .spec.ordinals

#### Description

StatefulSetOrdinals describes the policy used for replica ordinal assignment in this StatefulSet.

#### Type

**object**

Property	Type	Description
<b>start</b>	<b>integer</b>	start is the number representing the first replica's index. It may be used to number replicas from an alternate index (eg: 1-indexed) over the default 0-indexed names, or to orchestrate progressive movement of replicas from one StatefulSet to another. If set, replica indices will be in the range: [.spec.ordinals.start, .spec.ordinals.start + .spec.replicas). If unset, defaults to 0. Replica indices will be in the range: [0, .spec.replicas).

### 7.6.1.3. .spec.persistentVolumeClaimRetentionPolicy

#### Description

StatefulSetPersistentVolumeClaimRetentionPolicy describes the policy used for PVCs created from the StatefulSet VolumeClaimTemplates.

#### Type

**object**

Property	Type	Description
<b>whenDeleted</b>	<b>string</b>	WhenDeleted specifies what happens to PVCs created from StatefulSet VolumeClaimTemplates when the StatefulSet is deleted. The default policy of <b>Retain</b> causes PVCs to not be affected by StatefulSet deletion. The <b>Delete</b> policy causes those PVCs to be deleted.
<b>whenScaled</b>	<b>string</b>	WhenScaled specifies what happens to PVCs created from StatefulSet VolumeClaimTemplates when the StatefulSet is scaled down. The default policy of <b>Retain</b> causes PVCs to not be affected by a scaledown. The <b>Delete</b> policy causes the associated PVCs for any excess pods above the replica count to be deleted.

### 7.6.1.4. .spec.updateStrategy

#### Description

StatefulSetUpdateStrategy indicates the strategy that the StatefulSet controller will use to perform updates. It includes any additional parameters necessary to perform the update for the indicated strategy.

#### Type

**object**

Property	Type	Description
<b>rollingUpdate</b>	<b>object</b>	RollingUpdateStatefulSetStrategy is used to communicate parameter for RollingUpdateStatefulSetStrategyType.
<b>type</b>	<b>string</b>	<p>Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.</p> <p>Possible enum values: -  <b>"OnDelete"</b> triggers the legacy behavior. Version tracking and ordered rolling restarts are disabled. Pods are recreated from the StatefulSetSpec when they are manually deleted. When a scale operation is performed with this strategy, specification version indicated by the StatefulSet's currentRevision. -  <b>"RollingUpdate"</b> indicates that update will be applied to all Pods in the StatefulSet with respect to the StatefulSet ordering constraints. When a scale operation is performed with this strategy, new Pods will be created from the specification version indicated by the StatefulSet's updateRevision.</p>

### 7.6.1.5. .spec.updateStrategy.rollingUpdate

#### Description

RollingUpdateStatefulSetStrategy is used to communicate parameter for RollingUpdateStatefulSetStrategyType.

#### Type

**object**

Property	Type	Description
<b>maxUnavailable</b>	<b>IntOrString</b>	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding up. This can not be 0. Defaults to 1. This field is alpha-level and is only honored by servers that enable the MaxUnavailableStatefulSet feature. The field applies to all pods in the range 0 to Replicas-1. That means if there is any unavailable pod in the range 0 to Replicas-1, it will be counted towards MaxUnavailable.
<b>partition</b>	<b>integer</b>	Partition indicates the ordinal at which the StatefulSet should be partitioned for updates. During a rolling update, all pods from ordinal Replicas-1 to Partition are updated. All pods from ordinal Partition-1 to 0 remain untouched. This is helpful in being able to do a canary based deployment. The default value is 0.

### 7.6.1.6. .status

#### Description

StatefulSetStatus represents the current state of a StatefulSet.

#### Type

**object**

#### Required

- **replicas**

Property	Type	Description
<b>availableReplicas</b>	<b>integer</b>	Total number of available pods (ready for at least minReadySeconds) targeted by this statefulset.



Property	Type	Description
<b>collisionCount</b>	<b>integer</b>	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
<b>conditions</b>	<b>array</b>	Represents the latest available observations of a statefulset's current state.
<b>conditions[]</b>	<b>object</b>	StatefulSetCondition describes the state of a statefulset at a certain point.
<b>currentReplicas</b>	<b>integer</b>	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
<b>currentRevision</b>	<b>string</b>	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
<b>observedGeneration</b>	<b>integer</b>	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
<b>readyReplicas</b>	<b>integer</b>	readyReplicas is the number of pods created for this StatefulSet with a Ready Condition.
<b>replicas</b>	<b>integer</b>	replicas is the number of Pods created by the StatefulSet controller.
<b>updateRevision</b>	<b>string</b>	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

Property	Type	Description
<b>updatedReplicas</b>	<b>integer</b>	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

### 7.6.1.7. .status.conditions

#### Description

Represents the latest available observations of a statefulset's current state.

#### Type

**array**

### 7.6.1.8. .status.conditions[]

#### Description

StatefulSetCondition describes the state of a statefulset at a certain point.

#### Type

**object**

#### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	A human readable message indicating details about the transition.
<b>reason</b>	<b>string</b>	The reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of statefulset condition.

## 7.6.2. API endpoints

The following API endpoints are available:

- **/apis/apps/v1/statefulsets**
  - **GET**: list or watch objects of kind StatefulSet
- **/apis/apps/v1/watch/statefulsets**
  - **GET**: watch individual changes to a list of StatefulSet. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/statefulsets**
  - **DELETE**: delete collection of StatefulSet
  - **GET**: list or watch objects of kind StatefulSet
  - **POST**: create a StatefulSet
- **/apis/apps/v1/watch/namespaces/{namespace}/statefulsets**
  - **GET**: watch individual changes to a list of StatefulSet. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}**
  - **DELETE**: delete a StatefulSet
  - **GET**: read the specified StatefulSet
  - **PATCH**: partially update the specified StatefulSet
  - **PUT**: replace the specified StatefulSet
- **/apis/apps/v1/watch/namespaces/{namespace}/statefulsets/{name}**
  - **GET**: watch changes to an object of kind StatefulSet. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status**
  - **GET**: read status of the specified StatefulSet
  - **PATCH**: partially update status of the specified StatefulSet
  - **PUT**: replace status of the specified StatefulSet

### 7.6.2.1. /apis/apps/v1/statefulsets

Table 7.156. Global query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind StatefulSet

Table 7.157. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">StatefulSetList</a> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 7.6.2.2. /apis/apps/v1/watch/statefulsets

Table 7.158. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of StatefulSet. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.159. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.6.2.3. /apis/apps/v1/namespaces/{namespace}/statefulsets

Table 7.160. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.161. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of StatefulSet

Table 7.162. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 7.163. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.164. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind StatefulSet

Table 7.165. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 7.166. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">StatefulSetList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a StatefulSet

**Table 7.167. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.168. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>StatefulSet</b> schema	

**Table 7.169. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>StatefulSet</b> schema

HTTP code	Response body
201 - Created	<b>StatefulSet</b> schema
202 - Accepted	<b>StatefulSet</b> schema
401 - Unauthorized	Empty

#### 7.6.2.4. /apis/apps/v1/watch/namespaces/{namespace}/statefulsets

Table 7.170. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.171. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of StatefulSet. deprecated: use the 'watch' parameter with a list operation instead.

Table 7.172. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.6.2.5. /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}

Table 7.173. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StatefulSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.174. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a StatefulSet

Table 7.175. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 7.176. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 7.177. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified StatefulSet

Table 7.178. HTTP responses

HTTP code	Response body
200 - OK	<b>StatefulSet</b> schema



HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified StatefulSet

**Table 7.179. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.180. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.181. HTTP responses

HTTP code	Response body
200 - OK	<b>StatefulSet</b> schema
201 - Created	<b>StatefulSet</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified StatefulSet

Table 7.182. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 7.183. Body parameters

Parameter	Type	Description
<b>body</b>	<b>StatefulSet</b> schema	

Table 7.184. HTTP responses

HTTP code	Response body
200 - OK	<b>StatefulSet</b> schema
201 - Created	<b>StatefulSet</b> schema
401 - Unauthorized	Empty

#### 7.6.2.6. /apis/apps/v1/watch/namespaces/{namespace}/statefulsets/{name}

Table 7.185. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StatefulSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.186. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind StatefulSet. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 7.187. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 7.6.2.7. /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status

Table 7.188. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StatefulSet
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 7.189. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified StatefulSet

Table 7.190. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">StatefulSet</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified StatefulSet

Table 7.191. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 7.192. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 7.193. HTTP responses



HTTP code	Response body
200 - OK	<b>StatefulSet</b> schema
201 - Created	<b>StatefulSet</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified StatefulSet

**Table 7.194. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 7.195. Body parameters**

Parameter	Type	Description
<b>body</b>	<a href="#">StatefulSet</a> schema	

Table 7.196. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">StatefulSet</a> schema
201 - Created	<a href="#">StatefulSet</a> schema
401 - Unauthorized	Empty

## CHAPTER 8. AUTHENTICATION APIS

### 8.1. AUTHENTICATION APIS

#### 8.1.1. TokenRequest [authentication.k8s.io/v1]

##### Description

TokenRequest requests a token for a given service account.

##### Type

**object**

#### 8.1.2. TokenReview [authentication.k8s.io/v1]

##### Description

TokenReview attempts to authenticate a token to a known user. Note: TokenReview requests may be cached by the webhook token authenticator plugin in the kube-apiserver.

##### Type

**object**

### 8.2. TOKENREQUEST [AUTHENTICATION.K8S.IO/V1]

##### Description

TokenRequest requests a token for a given service account.

##### Type

**object**

##### Required

- **spec**

#### 8.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	TokenRequestSpec contains client provided parameters of a token request.
<b>status</b>	<b>object</b>	TokenRequestStatus is the result of a token request.

### 8.2.1.1. .spec

#### Description

TokenRequestSpec contains client provided parameters of a token request.

#### Type

**object**

#### Required

- **audiences**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>audiences</b>	<b>array (string)</b>	Audiences are the intended audiences of the token. A recipient of a token must identify themselves with an identifier in the list of audiences of the token, and otherwise should reject the token. A token issued for multiple audiences may be used to authenticate against any of the audiences listed but implies a high degree of trust between the target audiences.
<b>boundObjectRef</b>	<b>object</b>	BoundObjectReference is a reference to an object that a token is bound to.
<b>expirationSeconds</b>	<b>integer</b>	ExpirationSeconds is the requested duration of validity of the request. The token issuer may return a token with a different validity duration so a client needs to check the 'expiration' field in a response.

### 8.2.1.2. .spec.boundObjectRef

#### Description

BoundObjectReference is a reference to an object that a token is bound to.

#### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>kind</b>	<b>string</b>	Kind of the referent. Valid kinds are 'Pod' and 'Secret'.
<b>name</b>	<b>string</b>	Name of the referent.
<b>uid</b>	<b>string</b>	UID of the referent.

### 8.2.1.3. .status

#### Description

TokenRequestStatus is the result of a token request.

## Type

**object**

## Required

- **token**
- **expirationTimestamp**

Property	Type	Description
<b>expirationTimestamp</b>	<b>Time</b>	ExpirationTimestamp is the time of expiration of the returned token.
<b>token</b>	<b>string</b>	Token is the opaque bearer token.

## 8.2.2. API endpoints

The following API endpoints are available:

- **/api/v1/namespaces/{namespace}/serviceaccounts/{name}/token**
  - **POST**: create token of a ServiceAccount

### 8.2.2.1. /api/v1/namespaces/{namespace}/serviceaccounts/{name}/token

Table 8.1. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the TokenRequest
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 8.2. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create token of a ServiceAccount

**Table 8.3. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>TokenRequest</b> schema	

**Table 8.4. HTTP responses**

HTTP code	Response body
200 - OK	<b>TokenRequest</b> schema

HTTP code	Reponse body
201 - Created	<b>TokenRequest</b> schema
202 - Accepted	<b>TokenRequest</b> schema
401 - Unauthorized	Empty

## 8.3. TOKENREVIEW [AUTHENTICATION.K8S.IO/V1]

### Description

TokenReview attempts to authenticate a token to a known user. Note: TokenReview requests may be cached by the webhook token authenticator plugin in the kube-apiserver.

### Type

**object**

### Required

- **spec**

### 8.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>



Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	TokenReviewSpec is a description of the token authentication request.
<b>status</b>	<b>object</b>	TokenReviewStatus is the result of the token authentication request.

### 8.3.1.1. .spec

#### Description

TokenReviewSpec is a description of the token authentication request.

#### Type

**object**

Property	Type	Description
<b>audiences</b>	<b>array (string)</b>	Audiences is a list of the identifiers that the resource server presented with the token identifies as. Audience-aware token authenticators will verify that the token was intended for at least one of the audiences in this list. If no audiences are provided, the audience will default to the audience of the Kubernetes apiserver.
<b>token</b>	<b>string</b>	Token is the opaque bearer token.

### 8.3.1.2. .status

#### Description

TokenReviewStatus is the result of the token authentication request.

#### Type

**object**

Property	Type	Description
<b>audiences</b>	<b>array (string)</b>	Audiences are audience identifiers chosen by the authenticator that are compatible with both the TokenReview and token. An identifier is any identifier in the intersection of the TokenReviewSpec audiences and the token's audiences. A client of the TokenReview API that sets the spec.audiences field should validate that a compatible audience identifier is returned in the status.audiences field to ensure that the TokenReview server is audience aware. If a TokenReview returns an empty status.audience field where status.authenticated is "true", the token is valid against the audience of the Kubernetes API server.
<b>authenticated</b>	<b>boolean</b>	Authenticated indicates that the token was associated with a known user.
<b>error</b>	<b>string</b>	Error indicates that the token couldn't be checked
<b>user</b>	<b>object</b>	UserInfo holds the information about the user needed to implement the user.Info interface.

### 8.3.1.3. .status.user

#### Description

UserInfo holds the information about the user needed to implement the user.Info interface.

#### Type

**object**

Property	Type	Description
<b>extra</b>	<b>object</b>	Any additional information provided by the authenticator.
<b>extra{}</b>	<b>array (string)</b>	

Property	Type	Description
<b>groups</b>	<b>array (string)</b>	The names of groups this user is a part of.
<b>uid</b>	<b>string</b>	A unique value that identifies this user across time. If this user is deleted and another user by the same name is added, they will have different UIDs.
<b>username</b>	<b>string</b>	The name that uniquely identifies this user among all active users.

#### 8.3.1.4. .status.user.extra

##### Description

Any additional information provided by the authenticator.

##### Type

**object**

### 8.3.2. API endpoints

The following API endpoints are available:

- **/apis/authentication.k8s.io/v1/tokenreviews**
  - **POST**: create a TokenReview

#### 8.3.2.1. /apis/authentication.k8s.io/v1/tokenreviews

Table 8.5. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create a TokenReview

**Table 8.6. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>TokenReview</b> schema	

**Table 8.7. HTTP responses**

HTTP code	Response body
200 - OK	<b>TokenReview</b> schema
201 - Created	<b>TokenReview</b> schema
202 - Accepted	<b>TokenReview</b> schema
401 - Unauthorized	Empty

## CHAPTER 9. AUTHORIZATION APIS

### 9.1. AUTHORIZATION APIS

#### 9.1.1. LocalSubjectAccessReview [authorization.k8s.io/v1]

##### Description

LocalSubjectAccessReview checks whether or not a user or group can perform an action in a given namespace. Having a namespace scoped resource makes it much easier to grant namespace scoped policy that includes permissions checking.

##### Type

**object**

#### 9.1.2. SelfSubjectAccessReview [authorization.k8s.io/v1]

##### Description

SelfSubjectAccessReview checks whether or the current user can perform an action. Not filling in a spec.namespace means "in all namespaces". Self is a special case, because users should always be able to check whether they can perform an action

##### Type

**object**

#### 9.1.3. SelfSubjectRulesReview [authorization.k8s.io/v1]

##### Description

SelfSubjectRulesReview enumerates the set of actions the current user can perform within a namespace. The returned list of actions may be incomplete depending on the server's authorization mode, and any errors experienced during the evaluation. SelfSubjectRulesReview should be used by UIs to show/hide actions, or to quickly let an end user reason about their permissions. It should NOT Be used by external systems to drive authorization decisions as this raises confused deputy, cache lifetime/revocation, and correctness concerns. SubjectAccessReview, and LocalAccessReview are the correct way to defer authorization decisions to the API server.

##### Type

**object**

#### 9.1.4. SubjectAccessReview [authorization.k8s.io/v1]

##### Description

SubjectAccessReview checks whether or not a user or group can perform an action.

##### Type

**object**

### 9.2. LOCALSUBJECTACCESSREVIEW [AUTHORIZATION.K8S.IO/V1]

##### Description

LocalSubjectAccessReview checks whether or not a user or group can perform an action in a given namespace. Having a namespace scoped resource makes it much easier to grant namespace scoped policy that includes permissions checking.

Type

**object**

Required

- **spec**

### 9.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	SubjectAccessReviewSpec is a description of the access request. Exactly one of ResourceAuthorizationAttributes and NonResourceAuthorizationAttributes must be set
<b>status</b>	<b>object</b>	SubjectAccessReviewStatus

#### 9.2.1.1. .spec

**Description**

SubjectAccessReviewSpec is a description of the access request. Exactly one of ResourceAuthorizationAttributes and NonResourceAuthorizationAttributes must be set

**Type**

**object**

Property	Type	Description
<b>extra</b>	<b>object</b>	Extra corresponds to the user.Info.GetExtra() method from the authenticator. Since that is input to the authorizer it needs a reflection here.
<b>extra{}</b>	<b>array (string)</b>	
<b>groups</b>	<b>array (string)</b>	Groups is the groups you're testing for.
<b>nonResourceAttributes</b>	<b>object</b>	NonResourceAttributes includes the authorization attributes available for non-resource requests to the Authorizer interface
<b>resourceAttributes</b>	<b>object</b>	ResourceAttributes includes the authorization attributes available for resource requests to the Authorizer interface
<b>uid</b>	<b>string</b>	UID information about the requesting user.
<b>user</b>	<b>string</b>	User is the user you're testing for. If you specify "User" but not "Groups", then is it interpreted as "What if User were not a member of any groups"

**9.2.1.2. .spec.extra****Description**

Extra corresponds to the user.Info.GetExtra() method from the authenticator. Since that is input to the authorizer it needs a reflection here.

**Type**

**object**

**9.2.1.3. .spec.nonResourceAttributes****Description**

NonResourceAttributes includes the authorization attributes available for non-resource requests to the Authorizer interface

### Type

**object**

Property	Type	Description
<b>path</b>	<b>string</b>	Path is the URL path of the request
<b>verb</b>	<b>string</b>	Verb is the standard HTTP verb

### 9.2.1.4. .spec.resourceAttributes

#### Description

ResourceAttributes includes the authorization attributes available for resource requests to the Authorizer interface

### Type

**object**

Property	Type	Description
<b>group</b>	<b>string</b>	Group is the API Group of the Resource. "*" means all.
<b>name</b>	<b>string</b>	Name is the name of the resource being requested for a "get" or deleted for a "delete". "" (empty) means all.
<b>namespace</b>	<b>string</b>	Namespace is the namespace of the action being requested. Currently, there is no distinction between no namespace and all namespaces "" (empty) is defaulted for LocalSubjectAccessReviews "" (empty) is empty for cluster-scoped resources "" (empty) means "all" for namespace scoped resources from a SubjectAccessReview or SelfSubjectAccessReview
<b>resource</b>	<b>string</b>	Resource is one of the existing resource types. "*" means all.
<b>subresource</b>	<b>string</b>	Subresource is one of the existing resource types. "" means none.



Property	Type	Description
<b>verb</b>	<b>string</b>	Verb is a kubernetes resource API verb, like: get, list, watch, create, update, delete, proxy. "*" means all.
<b>version</b>	<b>string</b>	Version is the API Version of the Resource. "*" means all.

### 9.2.1.5. .status

#### Description

SubjectAccessReviewStatus

#### Type

**object**

#### Required

- **allowed**

Property	Type	Description
<b>allowed</b>	<b>boolean</b>	Allowed is required. True if the action would be allowed, false otherwise.
<b>denied</b>	<b>boolean</b>	Denied is optional. True if the action would be denied, otherwise false. If both allowed is false and denied is false, then the authorizer has no opinion on whether to authorize the action. Denied may not be true if Allowed is true.
<b>evaluationError</b>	<b>string</b>	EvaluationError is an indication that some error occurred during the authorization check. It is entirely possible to get an error and be able to continue determine authorization status in spite of it. For instance, RBAC can be missing a role, but enough roles are still present and bound to reason about the request.
<b>reason</b>	<b>string</b>	Reason is optional. It indicates why a request was allowed or denied.

## 9.2.2. API endpoints

The following API endpoints are available:

- **/apis/authorization.k8s.io/v1/namespaces/{namespace}/localsubjectaccessreviews**
  - **POST**: create a LocalSubjectAccessReview

### 9.2.2.1. /apis/authorization.k8s.io/v1/namespaces/{namespace}/localsubjectaccessreviews

Table 9.1. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 9.2. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**POST**

Description

create a LocalSubjectAccessReview

Table 9.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>LocalSubjectAccessReview</b> schema	

Table 9.4. HTTP responses

HTTP code	Reponse body
200 - OK	<b>LocalSubjectAccessReview</b> schema
201 - Created	<b>LocalSubjectAccessReview</b> schema
202 - Accepted	<b>LocalSubjectAccessReview</b> schema
401 - Unauthorized	Empty

## 9.3. SELFSUBJECTACCESSREVIEW [AUTHORIZATION.K8S.IO/V1]

Description

SelfSubjectAccessReview checks whether or the current user can perform an action. Not filling in a spec.namespace means "in all namespaces". Self is a special case, because users should always be able to check whether they can perform an action

Type

**object**

Required

- **spec**

### 9.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	SelfSubjectAccessReviewSpec is a description of the access request. Exactly one of ResourceAuthorizationAttributes and NonResourceAuthorizationAttributes must be set
<b>status</b>	<b>object</b>	SubjectAccessReviewStatus

### 9.3.1.1. .spec

#### Description

SelfSubjectAccessReviewSpec is a description of the access request. Exactly one of ResourceAuthorizationAttributes and NonResourceAuthorizationAttributes must be set

#### Type

**object**

Property	Type	Description
<b>nonResourceAttributes</b>	<b>object</b>	NonResourceAttributes includes the authorization attributes available for non-resource requests to the Authorizer interface
<b>resourceAttributes</b>	<b>object</b>	ResourceAttributes includes the authorization attributes available for resource requests to the Authorizer interface

### 9.3.1.2. .spec.nonResourceAttributes

#### Description

NonResourceAttributes includes the authorization attributes available for non-resource requests to the Authorizer interface

#### Type

**object**

Property	Type	Description
<b>path</b>	<b>string</b>	Path is the URL path of the request
<b>verb</b>	<b>string</b>	Verb is the standard HTTP verb

### 9.3.1.3. .spec.resourceAttributes

#### Description

ResourceAttributes includes the authorization attributes available for resource requests to the Authorizer interface

#### Type

**object**

Property	Type	Description
<b>group</b>	<b>string</b>	Group is the API Group of the Resource. "*" means all.
<b>name</b>	<b>string</b>	Name is the name of the resource being requested for a "get" or deleted for a "delete". "" (empty) means all.

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace is the namespace of the action being requested. Currently, there is no distinction between no namespace and all namespaces "" (empty) is defaulted for LocalSubjectAccessReviews "" (empty) is empty for cluster-scoped resources "" (empty) means "all" for namespace scoped resources from a SubjectAccessReview or SelfSubjectAccessReview
<b>resource</b>	<b>string</b>	Resource is one of the existing resource types. "*" means all.
<b>subresource</b>	<b>string</b>	Subresource is one of the existing resource types. "" means none.
<b>verb</b>	<b>string</b>	Verb is a kubernetes resource API verb, like: get, list, watch, create, update, delete, proxy. "*" means all.
<b>version</b>	<b>string</b>	Version is the API Version of the Resource. "*" means all.

#### 9.3.1.4. .status

##### Description

SubjectAccessReviewStatus

##### Type

**object**

##### Required

- **allowed**

Property	Type	Description
<b>allowed</b>	<b>boolean</b>	Allowed is required. True if the action would be allowed, false otherwise.

Property	Type	Description
<b>denied</b>	<b>boolean</b>	Denied is optional. True if the action would be denied, otherwise false. If both allowed is false and denied is false, then the authorizer has no opinion on whether to authorize the action. Denied may not be true if Allowed is true.
<b>evaluationError</b>	<b>string</b>	EvaluationError is an indication that some error occurred during the authorization check. It is entirely possible to get an error and be able to continue determine authorization status in spite of it. For instance, RBAC can be missing a role, but enough roles are still present and bound to reason about the request.
<b>reason</b>	<b>string</b>	Reason is optional. It indicates why a request was allowed or denied.

### 9.3.2. API endpoints

The following API endpoints are available:

- **/apis/authorization.k8s.io/v1/selfsubjectaccessreviews**
  - **POST**: create a SelfSubjectAccessReview

#### 9.3.2.1. /apis/authorization.k8s.io/v1/selfsubjectaccessreviews

Table 9.5. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create a SelfSubjectAccessReview

**Table 9.6. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>SelfSubjectAccessReview</b> schema	

**Table 9.7. HTTP responses**

HTTP code	Response body
200 - OK	<b>SelfSubjectAccessReview</b> schema
201 - Created	<b>SelfSubjectAccessReview</b> schema



HTTP code	Response body
202 - Accepted	<b>SelfSubjectAccessReview</b> schema
401 - Unauthorized	Empty

## 9.4. SELFSUBJECTRULESREVIEW [AUTHORIZATION.K8S.IO/V1]

### Description

SelfSubjectRulesReview enumerates the set of actions the current user can perform within a namespace. The returned list of actions may be incomplete depending on the server's authorization mode, and any errors experienced during the evaluation. SelfSubjectRulesReview should be used by UIs to show/hide actions, or to quickly let an end user reason about their permissions. It should NOT Be used by external systems to drive authorization decisions as this raises confused deputy, cache lifetime/revocation, and correctness concerns. SubjectAccessReview, and LocalAccessReview are the correct way to defer authorization decisions to the API server.

### Type

**object**

### Required

- **spec**

### 9.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	SelfSubjectRulesReviewSpec defines the specification for SelfSubjectRulesReview.
<b>status</b>	<b>object</b>	SubjectRulesReviewStatus contains the result of a rules check. This check can be incomplete depending on the set of authorizers the server is configured with and any errors experienced during evaluation. Because authorization rules are additive, if a rule appears in a list it's safe to assume the subject has that permission, even if that list is incomplete.

#### 9.4.1.1. .spec

##### Description

SelfSubjectRulesReviewSpec defines the specification for SelfSubjectRulesReview.

##### Type

**object**

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace to evaluate rules for. Required.

#### 9.4.1.2. .status

##### Description

SubjectRulesReviewStatus contains the result of a rules check. This check can be incomplete depending on the set of authorizers the server is configured with and any errors experienced during evaluation. Because authorization rules are additive, if a rule appears in a list it's safe to assume the subject has that permission, even if that list is incomplete.

##### Type

**object**

Required

- **resourceRules**
- **nonResourceRules**
- **incomplete**

Property	Type	Description
<b>evaluationError</b>	<b>string</b>	EvaluationError can appear in combination with Rules. It indicates an error occurred during rule evaluation, such as an authorizer that doesn't support rule evaluation, and that ResourceRules and/or NonResourceRules may be incomplete.
<b>incomplete</b>	<b>boolean</b>	Incomplete is true when the rules returned by this call are incomplete. This is most commonly encountered when an authorizer, such as an external authorizer, doesn't support rules evaluation.
<b>nonResourceRules</b>	<b>array</b>	NonResourceRules is the list of actions the subject is allowed to perform on non-resources. The list ordering isn't significant, may contain duplicates, and possibly be incomplete.
<b>nonResourceRules[]</b>	<b>object</b>	NonResourceRule holds information that describes a rule for the non-resource
<b>resourceRules</b>	<b>array</b>	ResourceRules is the list of actions the subject is allowed to perform on resources. The list ordering isn't significant, may contain duplicates, and possibly be incomplete.
<b>resourceRules[]</b>	<b>object</b>	ResourceRule is the list of actions the subject is allowed to perform on resources. The list ordering isn't significant, may contain duplicates, and possibly be incomplete.

### 9.4.1.3. .status.nonResourceRules

#### Description

NonResourceRules is the list of actions the subject is allowed to perform on non-resources. The list ordering isn't significant, may contain duplicates, and possibly be incomplete.

#### Type

**array**

### 9.4.1.4. .status.nonResourceRules[]

#### Description

NonResourceRule holds information that describes a rule for the non-resource

#### Type

**object**

#### Required

- **verbs**

Property	Type	Description
<b>nonResourceURLs</b>	<b>array (string)</b>	NonResourceURLs is a set of partial urls that a user should have access to. <b>s are allowed, but only as the full, final step in the path.</b> "" means all.
<b>verbs</b>	<b>array (string)</b>	Verb is a list of kubernetes non-resource API verbs, like: get, post, put, delete, patch, head, options. "*" means all.

### 9.4.1.5. .status.resourceRules

#### Description

ResourceRules is the list of actions the subject is allowed to perform on resources. The list ordering isn't significant, may contain duplicates, and possibly be incomplete.

#### Type

**array**

### 9.4.1.6. .status.resourceRules[]

#### Description

ResourceRule is the list of actions the subject is allowed to perform on resources. The list ordering isn't significant, may contain duplicates, and possibly be incomplete.

#### Type

**object**

#### Required

- **verbs**

- **verbs**

Property	Type	Description
<b>apiGroups</b>	<b>array (string)</b>	APIGroups is the name of the APIGroup that contains the resources. If multiple API groups are specified, any action requested against one of the enumerated resources in any API group will be allowed. "*" means all.
<b>resourceNames</b>	<b>array (string)</b>	ResourceNames is an optional white list of names that the rule applies to. An empty set means that everything is allowed. "*" means all.
<b>resources</b>	<b>array (string)</b>	Resources is a list of resources this rule applies to. "" <b>means all in the specified apiGroups.</b> "/foo" represents the subresource 'foo' for all resources in the specified apiGroups.
<b>verbs</b>	<b>array (string)</b>	Verb is a list of kubernetes resource API verbs, like: get, list, watch, create, update, delete, proxy. "*" means all.

## 9.4.2. API endpoints

The following API endpoints are available:

- **/apis/authorization.k8s.io/v1/selfsubjectrulesreviews**
  - **POST**: create a SelfSubjectRulesReview

### 9.4.2.1. /apis/authorization.k8s.io/v1/selfsubjectrulesreviews

Table 9.8. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create a SelfSubjectRulesReview

**Table 9.9. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>SelfSubjectRulesReview</b> schema	

**Table 9.10. HTTP responses**

HTTP code	Response body
200 - OK	<b>SelfSubjectRulesReview</b> schema

HTTP code	Response body
201 - Created	<b>SelfSubjectRulesReview</b> schema
202 - Accepted	<b>SelfSubjectRulesReview</b> schema
401 - Unauthorized	Empty

## 9.5. SUBJECTACCESSREVIEW [AUTHORIZATION.K8S.IO/V1]

### Description

SubjectAccessReview checks whether or not a user or group can perform an action.

### Type

**object**

### Required

- **spec**

### 9.5.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard list metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	SubjectAccessReviewSpec is a description of the access request. Exactly one of ResourceAuthorizationAttributes and NonResourceAuthorizationAttributes must be set
<b>status</b>	<b>object</b>	SubjectAccessReviewStatus

### 9.5.1.1. .spec

#### Description

SubjectAccessReviewSpec is a description of the access request. Exactly one of ResourceAuthorizationAttributes and NonResourceAuthorizationAttributes must be set

#### Type

**object**

Property	Type	Description
<b>extra</b>	<b>object</b>	Extra corresponds to the user.Info.GetExtra() method from the authenticator. Since that is input to the authorizer it needs a reflection here.
<b>extra{}</b>	<b>array (string)</b>	
<b>groups</b>	<b>array (string)</b>	Groups is the groups you're testing for.
<b>nonResourceAttributes</b>	<b>object</b>	NonResourceAttributes includes the authorization attributes available for non-resource requests to the Authorizer interface



Property	Type	Description
<b>resourceAttributes</b>	<b>object</b>	ResourceAttributes includes the authorization attributes available for resource requests to the Authorizer interface
<b>uid</b>	<b>string</b>	UID information about the requesting user.
<b>user</b>	<b>string</b>	User is the user you're testing for. If you specify "User" but not "Groups", then is it interpreted as "What if User were not a member of any groups"

### 9.5.1.2. .spec.extra

#### Description

Extra corresponds to the `user.Info.GetExtra()` method from the authenticator. Since that is input to the authorizer it needs a reflection here.

#### Type

**object**

### 9.5.1.3. .spec.nonResourceAttributes

#### Description

NonResourceAttributes includes the authorization attributes available for non-resource requests to the Authorizer interface

#### Type

**object**

Property	Type	Description
<b>path</b>	<b>string</b>	Path is the URL path of the request
<b>verb</b>	<b>string</b>	Verb is the standard HTTP verb

### 9.5.1.4. .spec.resourceAttributes

#### Description

ResourceAttributes includes the authorization attributes available for resource requests to the Authorizer interface

#### Type

**object**

Property	Type	Description
<b>group</b>	<b>string</b>	Group is the API Group of the Resource. "*" means all.
<b>name</b>	<b>string</b>	Name is the name of the resource being requested for a "get" or deleted for a "delete". "" (empty) means all.
<b>namespace</b>	<b>string</b>	Namespace is the namespace of the action being requested. Currently, there is no distinction between no namespace and all namespaces "" (empty) is defaulted for LocalSubjectAccessReviews "" (empty) is empty for cluster-scoped resources "" (empty) means "all" for namespace scoped resources from a SubjectAccessReview or SelfSubjectAccessReview
<b>resource</b>	<b>string</b>	Resource is one of the existing resource types. "*" means all.
<b>subresource</b>	<b>string</b>	Subresource is one of the existing resource types. "" means none.
<b>verb</b>	<b>string</b>	Verb is a kubernetes resource API verb, like: get, list, watch, create, update, delete, proxy. "*" means all.
<b>version</b>	<b>string</b>	Version is the API Version of the Resource. "*" means all.

### 9.5.1.5. .status

#### Description

SubjectAccessReviewStatus

#### Type

**object**

#### Required

- **allowed**

Property	Type	Description
<b>allowed</b>	<b>boolean</b>	Allowed is required. True if the action would be allowed, false otherwise.
<b>denied</b>	<b>boolean</b>	Denied is optional. True if the action would be denied, otherwise false. If both allowed is false and denied is false, then the authorizer has no opinion on whether to authorize the action. Denied may not be true if Allowed is true.
<b>evaluationError</b>	<b>string</b>	EvaluationError is an indication that some error occurred during the authorization check. It is entirely possible to get an error and be able to continue determine authorization status in spite of it. For instance, RBAC can be missing a role, but enough roles are still present and bound to reason about the request.
<b>reason</b>	<b>string</b>	Reason is optional. It indicates why a request was allowed or denied.

## 9.5.2. API endpoints

The following API endpoints are available:

- **/apis/authorization.k8s.io/v1/subjectaccessreviews**
  - **POST**: create a SubjectAccessReview

### 9.5.2.1. /apis/authorization.k8s.io/v1/subjectaccessreviews

Table 9.11. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create a SubjectAccessReview

**Table 9.12. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>SubjectAccessReview</b> schema	

**Table 9.13. HTTP responses**

HTTP code	Response body
200 - OK	<b>SubjectAccessReview</b> schema
201 - Created	<b>SubjectAccessReview</b> schema

HTTP code	Reponse body
202 - Accepted	<b>SubjectAccessReview</b> schema
401 - Unauthorized	Empty

## CHAPTER 10. AUTOSCALING APIS

### 10.1. AUTOSCALING APIS

#### 10.1.1. Scale [autoscaling/v1]

##### Description

Scale represents a scaling request for a resource.

##### Type

**object**

#### 10.1.2. HorizontalPodAutoscaler [autoscaling/v2]

##### Description

HorizontalPodAutoscaler is the configuration for a horizontal pod autoscaler, which automatically manages the replica count of any resource implementing the scale subresource based on the metrics specified.

##### Type

**object**

### 10.2. HORIZONTALPODAUTOSCALER [AUTOSCALING/V2]

##### Description

HorizontalPodAutoscaler is the configuration for a horizontal pod autoscaler, which automatically manages the replica count of any resource implementing the scale subresource based on the metrics specified.

##### Type

**object**

#### 10.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	metadata is the standard object metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	HorizontalPodAutoscalerSpec describes the desired functionality of the HorizontalPodAutoscaler.
<b>status</b>	<b>object</b>	HorizontalPodAutoscalerStatus describes the current status of a horizontal pod autoscaler.

### 10.2.1.1. .spec

#### Description

HorizontalPodAutoscalerSpec describes the desired functionality of the HorizontalPodAutoscaler.

#### Type

**object**

#### Required

- **scaleTargetRef**
- **maxReplicas**

Property	Type	Description
<b>behavior</b>	<b>object</b>	HorizontalPodAutoscalerBehavior configures the scaling behavior of the target in both Up and Down directions (scaleUp and scaleDown fields respectively).

Property	Type	Description
<b>maxReplicas</b>	<b>integer</b>	maxReplicas is the upper limit for the number of replicas to which the autoscaler can scale up. It cannot be less than minReplicas.
<b>metrics</b>	<b>array</b>	metrics contains the specifications for which to use to calculate the desired replica count (the maximum replica count across all metrics will be used). The desired replica count is calculated multiplying the ratio between the target value and the current value by the current number of pods. Ergo, metrics used must decrease as the pod count is increased, and vice-versa. See the individual metric source types for more information about how each type of metric must respond. If not set, the default metric will be set to 80% average CPU utilization.
<b>metrics[]</b>	<b>object</b>	MetricSpec specifies how to scale based on a single metric (only <b>type</b> and one other matching field should be set at once).
<b>minReplicas</b>	<b>integer</b>	minReplicas is the lower limit for the number of replicas to which the autoscaler can scale down. It defaults to 1 pod. minReplicas is allowed to be 0 if the alpha feature gate HPAScaleToZero is enabled and at least one Object or External metric is configured. Scaling is active as long as at least one metric value is available.
<b>scaleTargetRef</b>	<b>object</b>	CrossVersionObjectReference contains enough information to let you identify the referred resource.

### 10.2.1.2. .spec.behavior

#### Description

HorizontalPodAutoscalerBehavior configures the scaling behavior of the target in both Up and Down directions (scaleUp and scaleDown fields respectively).



Type  
object

Property	Type	Description
<b>scaleDown</b>	<b>object</b>	HPAScalingRules configures the scaling behavior for one direction. These Rules are applied after calculating DesiredReplicas from metrics for the HPA. They can limit the scaling velocity by specifying scaling policies. They can prevent flapping by specifying the stabilization window, so that the number of replicas is not set instantly, instead, the safest value from the stabilization window is chosen.
<b>scaleUp</b>	<b>object</b>	HPAScalingRules configures the scaling behavior for one direction. These Rules are applied after calculating DesiredReplicas from metrics for the HPA. They can limit the scaling velocity by specifying scaling policies. They can prevent flapping by specifying the stabilization window, so that the number of replicas is not set instantly, instead, the safest value from the stabilization window is chosen.

### 10.2.1.3. .spec.behavior.scaleDown

#### Description

HPAScalingRules configures the scaling behavior for one direction. These Rules are applied after calculating DesiredReplicas from metrics for the HPA. They can limit the scaling velocity by specifying scaling policies. They can prevent flapping by specifying the stabilization window, so that the number of replicas is not set instantly, instead, the safest value from the stabilization window is chosen.

Type  
object

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>policies</b>	<b>array</b>	policies is a list of potential scaling polices which can be used during scaling. At least one policy must be specified, otherwise the HPAScalingRules will be discarded as invalid
<b>policies[]</b>	<b>object</b>	HPAScalingPolicy is a single policy which must hold true for a specified past interval.
<b>selectPolicy</b>	<b>string</b>	selectPolicy is used to specify which policy should be used. If not set, the default value Max is used.
<b>stabilizationWindowSeconds</b>	<b>integer</b>	stabilizationWindowSeconds is the number of seconds for which past recommendations should be considered while scaling up or scaling down. StabilizationWindowSeconds must be greater than or equal to zero and less than or equal to 3600 (one hour). If not set, use the default values: - For scale up: 0 (i.e. no stabilization is done). - For scale down: 300 (i.e. the stabilization window is 300 seconds long).

#### 10.2.1.4. .spec.behavior.scaleDown.policies

##### Description

policies is a list of potential scaling polices which can be used during scaling. At least one policy must be specified, otherwise the HPAScalingRules will be discarded as invalid

##### Type

**array**

#### 10.2.1.5. .spec.behavior.scaleDown.policies[]

##### Description

HPAScalingPolicy is a single policy which must hold true for a specified past interval.

##### Type

**object**

##### Required

- **type**
- **value**

- **periodSeconds**

Property	Type	Description
<b>periodSeconds</b>	<b>integer</b>	periodSeconds specifies the window of time for which the policy should hold true. PeriodSeconds must be greater than zero and less than or equal to 1800 (30 min).
<b>type</b>	<b>string</b>	type is used to specify the scaling policy.
<b>value</b>	<b>integer</b>	value contains the amount of change which is permitted by the policy. It must be greater than zero

#### 10.2.1.6. .spec.behavior.scaleUp

##### Description

HPAScalingRules configures the scaling behavior for one direction. These Rules are applied after calculating DesiredReplicas from metrics for the HPA. They can limit the scaling velocity by specifying scaling policies. They can prevent flapping by specifying the stabilization window, so that the number of replicas is not set instantly, instead, the safest value from the stabilization window is chosen.

##### Type

**object**

Property	Type	Description
<b>policies</b>	<b>array</b>	policies is a list of potential scaling polices which can be used during scaling. At least one policy must be specified, otherwise the HPAScalingRules will be discarded as invalid
<b>policies[]</b>	<b>object</b>	HPAScalingPolicy is a single policy which must hold true for a specified past interval.
<b>selectPolicy</b>	<b>string</b>	selectPolicy is used to specify which policy should be used. If not set, the default value Max is used.

Property	Type	Description
<b>stabilizationWindowSeconds</b>	<b>integer</b>	<p>stabilizationWindowSeconds is the number of seconds for which past recommendations should be considered while scaling up or scaling down.</p> <p>StabilizationWindowSeconds must be greater than or equal to zero and less than or equal to 3600 (one hour). If not set, use the default values: - For scale up: 0 (i.e. no stabilization is done). - For scale down: 300 (i.e. the stabilization window is 300 seconds long).</p>

### 10.2.1.7. .spec.behavior.scaleUp.policies

#### Description

policies is a list of potential scaling polices which can be used during scaling. At least one policy must be specified, otherwise the HPAScalingRules will be discarded as invalid

#### Type

**array**

### 10.2.1.8. .spec.behavior.scaleUp.policies[]

#### Description

HPAScalingPolicy is a single policy which must hold true for a specified past interval.

#### Type

**object**

#### Required

- **type**
- **value**
- **periodSeconds**

Property	Type	Description
<b>periodSeconds</b>	<b>integer</b>	<p>periodSeconds specifies the window of time for which the policy should hold true.</p> <p>PeriodSeconds must be greater than zero and less than or equal to 1800 (30 min).</p>

Property	Type	Description
<b>type</b>	<b>string</b>	type is used to specify the scaling policy.
<b>value</b>	<b>integer</b>	value contains the amount of change which is permitted by the policy. It must be greater than zero

### 10.2.1.9. .spec.metrics

#### Description

metrics contains the specifications for which to use to calculate the desired replica count (the maximum replica count across all metrics will be used). The desired replica count is calculated multiplying the ratio between the target value and the current value by the current number of pods. Ergo, metrics used must decrease as the pod count is increased, and vice-versa. See the individual metric source types for more information about how each type of metric must respond. If not set, the default metric will be set to 80% average CPU utilization.

#### Type

**array**

### 10.2.1.10. .spec.metrics[]

#### Description

MetricSpec specifies how to scale based on a single metric (only **type** and one other matching field should be set at once).

#### Type

**object**

#### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>containerResource</b>	<b>object</b>	ContainerResourceMetricSource indicates how to scale on a resource metric known to Kubernetes, as specified in requests and limits, describing each pod in the current scale target (e.g. CPU or memory). The values will be averaged together before being compared to the target. Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source. Only one "target" type should be set.
<b>external</b>	<b>object</b>	ExternalMetricSource indicates how to scale on a metric not associated with any Kubernetes object (for example length of queue in cloud messaging service, or QPS from loadbalancer running outside of cluster).
<b>object</b>	<b>object</b>	ObjectMetricSource indicates how to scale on a metric describing a kubernetes object (for example, hits-per-second on an Ingress object).
<b>pods</b>	<b>object</b>	PodsMetricSource indicates how to scale on a metric describing each pod in the current scale target (for example, transactions-processed-per-second). The values will be averaged together before being compared to the target value.

Property	Type	Description
<b>resource</b>	<b>object</b>	ResourceMetricSource indicates how to scale on a resource metric known to Kubernetes, as specified in requests and limits, describing each pod in the current scale target (e.g. CPU or memory). The values will be averaged together before being compared to the target. Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source. Only one "target" type should be set.
<b>type</b>	<b>string</b>	type is the type of metric source. It should be one of "ContainerResource", "External", "Object", "Pods" or "Resource", each mapping to a matching field in the object. Note: "ContainerResource" type is available on when the feature-gate HPAContainerMetrics is enabled

### 10.2.1.11. .spec.metrics[].containerResource

#### Description

ContainerResourceMetricSource indicates how to scale on a resource metric known to Kubernetes, as specified in requests and limits, describing each pod in the current scale target (e.g. CPU or memory). The values will be averaged together before being compared to the target. Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source. Only one "target" type should be set.

#### Type

**object**

#### Required

- **name**
- **target**
- **container**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>container</b>	<b>string</b>	container is the name of the container in the pods of the scaling target
<b>name</b>	<b>string</b>	name is the name of the resource in question.
<b>target</b>	<b>object</b>	MetricTarget defines the target value, average value, or average utilization of a specific metric

### 10.2.1.12. .spec.metrics[].containerResource.target

#### Description

MetricTarget defines the target value, average value, or average utilization of a specific metric

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	averageUtilization is the target value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods. Currently only valid for Resource metric source type
<b>averageValue</b>	<b>Quantity</b>	averageValue is the target value of the average of the metric across all relevant pods (as a quantity)
<b>type</b>	<b>string</b>	type represents whether the metric type is Utilization, Value, or AverageValue
<b>value</b>	<b>Quantity</b>	value is the target value of the metric (as a quantity).



### 10.2.1.13. .spec.metrics[].external

#### Description

ExternalMetricSource indicates how to scale on a metric not associated with any Kubernetes object (for example length of queue in cloud messaging service, or QPS from loadbalancer running outside of cluster).

#### Type

**object**

#### Required

- **metric**
- **target**

Property	Type	Description
<b>metric</b>	<b>object</b>	MetricIdentifier defines the name and optionally selector for a metric
<b>target</b>	<b>object</b>	MetricTarget defines the target value, average value, or average utilization of a specific metric

### 10.2.1.14. .spec.metrics[].external.metric

#### Description

MetricIdentifier defines the name and optionally selector for a metric

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the given metric
<b>selector</b>	<b>LabelSelector</b>	selector is the string-encoded form of a standard kubernetes label selector for the given metric. When set, it is passed as an additional parameter to the metrics server for more specific metrics scoping. When unset, just the metricName will be used to gather metrics.

### 10.2.1.15. .spec.metrics[].external.target

#### Description

MetricTarget defines the target value, average value, or average utilization of a specific metric

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	averageUtilization is the target value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods. Currently only valid for Resource metric source type
<b>averageValue</b>	<b>Quantity</b>	averageValue is the target value of the average of the metric across all relevant pods (as a quantity)
<b>type</b>	<b>string</b>	type represents whether the metric type is Utilization, Value, or AverageValue
<b>value</b>	<b>Quantity</b>	value is the target value of the metric (as a quantity).

### 10.2.1.16. .spec.metrics[].object

#### Description

ObjectMetricSource indicates how to scale on a metric describing a kubernetes object (for example, hits-per-second on an Ingress object).

#### Type

**object**

#### Required

- **describedObject**
- **target**
- **metric**

Property	Type	Description
<b>describedObject</b>	<b>object</b>	CrossVersionObjectReference contains enough information to let you identify the referred resource.
<b>metric</b>	<b>object</b>	MetricIdentifier defines the name and optionally selector for a metric
<b>target</b>	<b>object</b>	MetricTarget defines the target value, average value, or average utilization of a specific metric

### 10.2.1.17. .spec.metrics[].object.describedObject

#### Description

CrossVersionObjectReference contains enough information to let you identify the referred resource.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	apiVersion is the API version of the referent
<b>kind</b>	<b>string</b>	kind is the kind of the referent; More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	name is the name of the referent; More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

### 10.2.1.18. .spec.metrics[].object.metric

#### Description

MetricIdentifier defines the name and optionally selector for a metric

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the given metric
<b>selector</b>	<a href="#">LabelSelector</a>	selector is the string-encoded form of a standard kubernetes label selector for the given metric. When set, it is passed as an additional parameter to the metrics server for more specific metrics scoping. When unset, just the metricName will be used to gather metrics.

#### 10.2.1.19. .spec.metrics[].object.target

#### Description

MetricTarget defines the target value, average value, or average utilization of a specific metric

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	averageUtilization is the target value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods. Currently only valid for Resource metric source type
<b>averageValue</b>	<a href="#">Quantity</a>	averageValue is the target value of the average of the metric across all relevant pods (as a quantity)

Property	Type	Description
<b>type</b>	<b>string</b>	type represents whether the metric type is Utilization, Value, or AverageValue
<b>value</b>	<b>Quantity</b>	value is the target value of the metric (as a quantity).

### 10.2.1.20. .spec.metrics[].pods

#### Description

PodsMetricSource indicates how to scale on a metric describing each pod in the current scale target (for example, transactions-processed-per-second). The values will be averaged together before being compared to the target value.

#### Type

**object**

#### Required

- **metric**
- **target**

Property	Type	Description
<b>metric</b>	<b>object</b>	MetricIdentifier defines the name and optionally selector for a metric
<b>target</b>	<b>object</b>	MetricTarget defines the target value, average value, or average utilization of a specific metric

### 10.2.1.21. .spec.metrics[].pods.metric

#### Description

MetricIdentifier defines the name and optionally selector for a metric

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the given metric
<b>selector</b>	<b>LabelSelector</b>	selector is the string-encoded form of a standard kubernetes label selector for the given metric. When set, it is passed as an additional parameter to the metrics server for more specific metrics scoping. When unset, just the metricName will be used to gather metrics.

### 10.2.1.22. .spec.metrics[].pods.target

#### Description

MetricTarget defines the target value, average value, or average utilization of a specific metric

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	averageUtilization is the target value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods. Currently only valid for Resource metric source type
<b>averageValue</b>	<b>Quantity</b>	averageValue is the target value of the average of the metric across all relevant pods (as a quantity)
<b>type</b>	<b>string</b>	type represents whether the metric type is Utilization, Value, or AverageValue
<b>value</b>	<b>Quantity</b>	value is the target value of the metric (as a quantity).

### 10.2.1.23. .spec.metrics[].resource

#### Description

ResourceMetricSource indicates how to scale on a resource metric known to Kubernetes, as specified in requests and limits, describing each pod in the current scale target (e.g. CPU or memory). The values will be averaged together before being compared to the target. Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source. Only one "target" type should be set.

#### Type

**object**

#### Required

- **name**
- **target**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the resource in question.
<b>target</b>	<b>object</b>	MetricTarget defines the target value, average value, or average utilization of a specific metric

### 10.2.1.24. .spec.metrics[].resource.target

#### Description

MetricTarget defines the target value, average value, or average utilization of a specific metric

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	averageUtilization is the target value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods. Currently only valid for Resource metric source type

Property	Type	Description
<b>averageValue</b>	<b>Quantity</b>	averageValue is the target value of the average of the metric across all relevant pods (as a quantity)
<b>type</b>	<b>string</b>	type represents whether the metric type is Utilization, Value, or AverageValue
<b>value</b>	<b>Quantity</b>	value is the target value of the metric (as a quantity).

### 10.2.1.25. .spec.scaleTargetRef

#### Description

CrossVersionObjectReference contains enough information to let you identify the referred resource.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	apiVersion is the API version of the referent
<b>kind</b>	<b>string</b>	kind is the kind of the referent; More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	name is the name of the referent; More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

### 10.2.1.26. .status

#### Description

HorizontalPodAutoscalerStatus describes the current status of a horizontal pod autoscaler.



Type

object

Required

- `desiredReplicas`

Property	Type	Description
<code>conditions</code>	array	<code>conditions</code> is the set of conditions required for this autoscaler to scale its target, and indicates whether or not those conditions are met.
<code>conditions[]</code>	object	<code>HorizontalPodAutoscalerCondition</code> describes the state of a <code>HorizontalPodAutoscaler</code> at a certain point.
<code>currentMetrics</code>	array	<code>currentMetrics</code> is the last read state of the metrics used by this autoscaler.
<code>currentMetrics[]</code>	object	<code>MetricStatus</code> describes the last-read state of a single metric.
<code>currentReplicas</code>	integer	<code>currentReplicas</code> is current number of replicas of pods managed by this autoscaler, as last seen by the autoscaler.
<code>desiredReplicas</code>	integer	<code>desiredReplicas</code> is the desired number of replicas of pods managed by this autoscaler, as last calculated by the autoscaler.
<code>lastScaleTime</code>	Time	<code>lastScaleTime</code> is the last time the <code>HorizontalPodAutoscaler</code> scaled the number of pods, used by the autoscaler to control how often the number of pods is changed.
<code>observedGeneration</code>	integer	<code>observedGeneration</code> is the most recent generation observed by this autoscaler.

### 10.2.1.27. `.status.conditions`

Description

HorizontalPodAutoscalerStatus conditions is an array of HorizontalPodAutoscalerCondition objects.

conditions is the set of conditions required for this autoscaler to scale its target, and indicates whether or not those conditions are met.

#### Type

**array**

### 10.2.1.28. .status.conditions[]

#### Description

HorizontalPodAutoscalerCondition describes the state of a HorizontalPodAutoscaler at a certain point.

#### Type

**object**

#### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	lastTransitionTime is the last time the condition transitioned from one status to another
<b>message</b>	<b>string</b>	message is a human-readable explanation containing details about the transition
<b>reason</b>	<b>string</b>	reason is the reason for the condition's last transition.
<b>status</b>	<b>string</b>	status is the status of the condition (True, False, Unknown)
<b>type</b>	<b>string</b>	type describes the current condition

### 10.2.1.29. .status.currentMetrics

#### Description

currentMetrics is the last read state of the metrics used by this autoscaler.

#### Type

**array**

### 10.2.1.30. .status.currentMetrics[]

#### Description

MetricStatus describes the last-read state of a single metric.

Type

**object**

Required

- **type**

Property	Type	Description
<b>containerResource</b>	<b>object</b>	ContainerResourceMetricStatus indicates the current value of a resource metric known to Kubernetes, as specified in requests and limits, describing a single container in each pod in the current scale target (e.g. CPU or memory). Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source.
<b>external</b>	<b>object</b>	ExternalMetricStatus indicates the current value of a global metric not associated with any Kubernetes object.
<b>object</b>	<b>object</b>	ObjectMetricStatus indicates the current value of a metric describing a kubernetes object (for example, hits-per-second on an Ingress object).
<b>Pods</b>	<b>object</b>	PodsMetricStatus indicates the current value of a metric describing each pod in the current scale target (for example, transactions-processed-per-second).
<b>resource</b>	<b>object</b>	ResourceMetricStatus indicates the current value of a resource metric known to Kubernetes, as specified in requests and limits, describing each pod in the current scale target (e.g. CPU or memory). Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source.

Property	Type	Description
<b>type</b>	<b>string</b>	type is the type of metric source. It will be one of "ContainerResource", "External", "Object", "Pods" or "Resource", each corresponds to a matching field in the object. Note: "ContainerResource" type is available on when the feature-gate HPAContainerMetrics is enabled

### 10.2.1.31. `.status.currentMetrics[].containerResource`

#### Description

ContainerResourceMetricStatus indicates the current value of a resource metric known to Kubernetes, as specified in requests and limits, describing a single container in each pod in the current scale target (e.g. CPU or memory). Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source.

#### Type

**object**

#### Required

- **name**
- **current**
- **container**

Property	Type	Description
<b>container</b>	<b>string</b>	container is the name of the container in the pods of the scaling target
<b>current</b>	<b>object</b>	MetricValueStatus holds the current value for a metric
<b>name</b>	<b>string</b>	name is the name of the resource in question.

### 10.2.1.32. `.status.currentMetrics[].containerResource.current`

#### Description

MetricValueStatus holds the current value for a metric

#### Type

**object**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	currentAverageUtilization is the current value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods.
<b>averageValue</b>	<b>Quantity</b>	averageValue is the current value of the average of the metric across all relevant pods (as a quantity)
<b>value</b>	<b>Quantity</b>	value is the current value of the metric (as a quantity).

**10.2.1.33. .status.currentMetrics[].external****Description**

ExternalMetricStatus indicates the current value of a global metric not associated with any Kubernetes object.

**Type****object****Required**

- **metric**
- **current**

Property	Type	Description
<b>current</b>	<b>object</b>	MetricValueStatus holds the current value for a metric
<b>metric</b>	<b>object</b>	MetricIdentifier defines the name and optionally selector for a metric

**10.2.1.34. .status.currentMetrics[].external.current****Description**

MetricValueStatus holds the current value for a metric

**Type****object**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	currentAverageUtilization is the current value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods.
<b>averageValue</b>	<b>Quantity</b>	averageValue is the current value of the average of the metric across all relevant pods (as a quantity)
<b>value</b>	<b>Quantity</b>	value is the current value of the metric (as a quantity).

### 10.2.1.35. .status.currentMetrics[.external.metric

#### Description

MetricIdentifier defines the name and optionally selector for a metric

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the given metric
<b>selector</b>	<b>LabelSelector</b>	selector is the string-encoded form of a standard kubernetes label selector for the given metric. When set, it is passed as an additional parameter to the metrics server for more specific metrics scoping. When unset, just the metricName will be used to gather metrics.

### 10.2.1.36. .status.currentMetrics[.object

#### Description

ObjectMetricStatus indicates the current value of a metric describing a kubernetes object (for example, hits-per-second on an Ingress object).

#### Type

**object**

## Required

- **metric**
- **current**
- **describedObject**

Property	Type	Description
<b>current</b>	<b>object</b>	MetricValueStatus holds the current value for a metric
<b>describedObject</b>	<b>object</b>	CrossVersionObjectReference contains enough information to let you identify the referred resource.
<b>metric</b>	<b>object</b>	MetricIdentifier defines the name and optionally selector for a metric

10.2.1.37. `.status.currentMetrics[].object.current`

## Description

MetricValueStatus holds the current value for a metric

## Type

**object**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	currentAverageUtilization is the current value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods.
<b>averageValue</b>	<b>Quantity</b>	averageValue is the current value of the average of the metric across all relevant pods (as a quantity)
<b>value</b>	<b>Quantity</b>	value is the current value of the metric (as a quantity).

10.2.1.38. `.status.currentMetrics[].object.describedObject`

**Description**

CrossVersionObjectReference contains enough information to let you identify the referred resource.

**Type**

**object**

**Required**

- **kind**
- **name**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	apiVersion is the API version of the referent
<b>kind</b>	<b>string</b>	kind is the kind of the referent; More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	name is the name of the referent; More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

**10.2.1.39. .status.currentMetrics[].object.metric****Description**

MetricIdentifier defines the name and optionally selector for a metric

**Type**

**object**

**Required**

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the given metric



Property	Type	Description
<b>selector</b>	<a href="#">LabelSelector</a>	selector is the string-encoded form of a standard kubernetes label selector for the given metric. When set, it is passed as an additional parameter to the metrics server for more specific metrics scoping. When unset, just the metricName will be used to gather metrics.

#### 10.2.1.40. .status.currentMetrics[].pods

##### Description

PodsMetricStatus indicates the current value of a metric describing each pod in the current scale target (for example, transactions-processed-per-second).

##### Type

**object**

##### Required

- **metric**
- **current**

Property	Type	Description
<b>current</b>	<b>object</b>	MetricValueStatus holds the current value for a metric
<b>metric</b>	<b>object</b>	MetricIdentifier defines the name and optionally selector for a metric

#### 10.2.1.41. .status.currentMetrics[].pods.current

##### Description

MetricValueStatus holds the current value for a metric

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	currentAverageUtilization is the current value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods.
<b>averageValue</b>	<b>Quantity</b>	averageValue is the current value of the average of the metric across all relevant pods (as a quantity)
<b>value</b>	<b>Quantity</b>	value is the current value of the metric (as a quantity).

#### 10.2.1.42. `.status.currentMetrics[].pods.metric`

##### Description

MetricIdentifier defines the name and optionally selector for a metric

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the given metric
<b>selector</b>	<b>LabelSelector</b>	selector is the string-encoded form of a standard kubernetes label selector for the given metric. When set, it is passed as an additional parameter to the metrics server for more specific metrics scoping. When unset, just the metricName will be used to gather metrics.

#### 10.2.1.43. `.status.currentMetrics[].resource`

##### Description

ResourceMetricStatus indicates the current value of a resource metric known to Kubernetes, as specified in requests and limits, describing each pod in the current scale target (e.g. CPU or

memory). Such metrics are built in to Kubernetes, and have special scaling options on top of those available to normal per-pod metrics using the "pods" source.

#### Type

**object**

#### Required

- **name**
- **current**

Property	Type	Description
<b>current</b>	<b>object</b>	MetricValueStatus holds the current value for a metric
<b>name</b>	<b>string</b>	name is the name of the resource in question.

#### 10.2.1.44. .status.currentMetrics[].resource.current

#### Description

MetricValueStatus holds the current value for a metric

#### Type

**object**

Property	Type	Description
<b>averageUtilization</b>	<b>integer</b>	currentAverageUtilization is the current value of the average of the resource metric across all relevant pods, represented as a percentage of the requested value of the resource for the pods.
<b>averageValue</b>	<b>Quantity</b>	averageValue is the current value of the average of the metric across all relevant pods (as a quantity)
<b>value</b>	<b>Quantity</b>	value is the current value of the metric (as a quantity).

#### 10.2.2. API endpoints

The following API endpoints are available:

- **/apis/autoscaling/v2/horizontalpodautoscalers**

- **GET**: list or watch objects of kind HorizontalPodAutoscaler
- **/apis/autoscaling/v2/watch/horizontalpodautoscalers**
  - **GET**: watch individual changes to a list of HorizontalPodAutoscaler. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers**
  - **DELETE**: delete collection of HorizontalPodAutoscaler
  - **GET**: list or watch objects of kind HorizontalPodAutoscaler
  - **POST**: create a HorizontalPodAutoscaler
- **/apis/autoscaling/v2/watch/namespaces/{namespace}/horizontalpodautoscalers**
  - **GET**: watch individual changes to a list of HorizontalPodAutoscaler. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers/{name}**
  - **DELETE**: delete a HorizontalPodAutoscaler
  - **GET**: read the specified HorizontalPodAutoscaler
  - **PATCH**: partially update the specified HorizontalPodAutoscaler
  - **PUT**: replace the specified HorizontalPodAutoscaler
- **/apis/autoscaling/v2/watch/namespaces/{namespace}/horizontalpodautoscalers/{name}**
  - **GET**: watch changes to an object of kind HorizontalPodAutoscaler. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers/{name}/status**
  - **GET**: read status of the specified HorizontalPodAutoscaler
  - **PATCH**: partially update status of the specified HorizontalPodAutoscaler
  - **PUT**: replace status of the specified HorizontalPodAutoscaler

#### 10.2.2.1. /apis/autoscaling/v2/horizontalpodautoscalers

Table 10.1. Global query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind HorizontalPodAutoscaler

**Table 10.2. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">HorizontalPodAutoscalerList</a> schema
401 - Unauthorized	Empty

### 10.2.2.2. /apis/autoscaling/v2/watch/horizontalpodautoscalers

Table 10.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.



Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue`</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of HorizontalPodAutoscaler. deprecated: use the 'watch' parameter with a list operation instead.

**Table 10.4. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 10.2.2.3. /apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers

Table 10.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of HorizontalPodAutoscaler

Table 10.7. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 10.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 10.9. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind HorizontalPodAutoscaler

Table 10.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 10.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">HorizontalPodAutoscalerList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a HorizontalPodAutoscaler

**Table 10.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 10.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>HorizontalPodAutoscaler</b> schema	

**Table 10.14. HTTP responses**

HTTP code	Response body
200 - OK	<b>HorizontalPodAutoscaler</b> schema

HTTP code	Response body
201 - Created	<a href="#">HorizontalPodAutoscaler</a> schema
202 - Accepted	<a href="#">HorizontalPodAutoscaler</a> schema
401 - Unauthorized	Empty

#### 10.2.2.4. /apis/autoscaling/v2/watch/namespaces/{namespace}/horizontalpodautoscalers

Table 10.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of HorizontalPodAutoscaler. deprecated: use the 'watch' parameter with a list operation instead.

Table 10.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 10.2.2.5. /apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers/{name}

Table 10.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the HorizontalPodAutoscaler
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a HorizontalPodAutoscaler

Table 10.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 10.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 10.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified HorizontalPodAutoscaler

Table 10.23. HTTP responses

HTTP code	Response body
200 - OK	<b>HorizontalPodAutoscaler</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified HorizontalPodAutoscaler

**Table 10.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 10.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 10.26. HTTP responses

HTTP code	Reponse body
200 - OK	<b>HorizontalPodAutoscaler</b> schema
201 - Created	<b>HorizontalPodAutoscaler</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified HorizontalPodAutoscaler

Table 10.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 10.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>HorizontalPodAutoscaler</b> schema	

Table 10.29. HTTP responses

HTTP code	Response body
200 - OK	<b>HorizontalPodAutoscaler</b> schema
201 - Created	<b>HorizontalPodAutoscaler</b> schema
401 - Unauthorized	Empty

### 10.2.2.6. /apis/autoscaling/v2/watch/namespaces/{namespace}/horizontalpodautoscalers/{

Table 10.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the HorizontalPodAutoscaler
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 10.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind HorizontalPodAutoscaler. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 10.32. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 10.2.2.7. /apis/autoscaling/v2/namespaces/{namespace}/horizontalpodautoscalers/{name}/

Table 10.33. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the HorizontalPodAutoscaler
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.34. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified HorizontalPodAutoscaler

Table 10.35. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">HorizontalPodAutoscaler</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified HorizontalPodAutoscaler

Table 10.36. Query parameters



Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 10.37. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 10.38. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">HorizontalPodAutoscaler</a> schema
201 - Created	<a href="#">HorizontalPodAutoscaler</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified HorizontalPodAutoscaler

**Table 10.39. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 10.40. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>HorizontalPodAutoscaler</b> schema	

Table 10.41. HTTP responses

HTTP code	Response body
200 - OK	<b>HorizontalPodAutoscaler</b> schema
201 - Created	<b>HorizontalPodAutoscaler</b> schema
401 - Unauthorized	Empty

## 10.3. SCALE [AUTOSCALING/V1]

### Description

Scale represents a scaling request for a resource.

### Type

**object**

### 10.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object metadata; More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a> .
<b>spec</b>	<b>object</b>	ScaleSpec describes the attributes of a scale subresource.
<b>status</b>	<b>object</b>	ScaleStatus represents the current status of a scale subresource.

### 10.3.1.1. .spec

#### Description

ScaleSpec describes the attributes of a scale subresource.

#### Type

**object**

Property	Type	Description
<b>replicas</b>	<b>integer</b>	replicas is the desired number of instances for the scaled object.

### 10.3.1.2. .status

#### Description

ScaleStatus represents the current status of a scale subresource.

#### Type

**object**

#### Required

- **replicas**

Property	Type	Description
<b>replicas</b>	<b>integer</b>	replicas is the actual number of observed instances of the scaled object.

Property	Type	Description
<b>selector</b>	<b>string</b>	selector is the label query over pods that should match the replicas count. This is same as the label selector but in the string format to avoid introspection by clients. The string will be in the same format as the query-param syntax. More info about label selectors: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/</a>

### 10.3.2. API endpoints

The following API endpoints are available:

- **/apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale**
  - **GET**: read scale of the specified Deployment
  - **PATCH**: partially update scale of the specified Deployment
  - **PUT**: replace scale of the specified Deployment
- **/apis/apps/v1/namespaces/{namespace}/replicasets/{name}/scale**
  - **GET**: read scale of the specified ReplicaSet
  - **PATCH**: partially update scale of the specified ReplicaSet
  - **PUT**: replace scale of the specified ReplicaSet
- **/apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/scale**
  - **GET**: read scale of the specified StatefulSet
  - **PATCH**: partially update scale of the specified StatefulSet
  - **PUT**: replace scale of the specified StatefulSet
- **/api/v1/namespaces/{namespace}/replicationcontrollers/{name}/scale**
  - **GET**: read scale of the specified ReplicationController
  - **PATCH**: partially update scale of the specified ReplicationController
  - **PUT**: replace scale of the specified ReplicationController

#### 10.3.2.1. /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale

Table 10.42. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Scale
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.43. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****GET****Description**

read scale of the specified Deployment

Table 10.44. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update scale of the specified Deployment

Table 10.45. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 10.46. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 10.47. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace scale of the specified Deployment

**Table 10.48. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 10.49. Body parameters**



Parameter	Type	Description
<b>body</b>	<b>Scale</b> schema	

Table 10.50. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema
401 - Unauthorized	Empty

### 10.3.2.2. /apis/apps/v1/namespaces/{namespace}/replicasets/{name}/scale

Table 10.51. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Scale
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.52. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read scale of the specified ReplicaSet

Table 10.53. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
401 - Unauthorized	Empty

#### HTTP method

**PATCH****Description**

partially update scale of the specified ReplicaSet

**Table 10.54. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

**Table 10.55. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 10.56. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace scale of the specified ReplicaSet

Table 10.57. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 10.58. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Scale</b> schema	

Table 10.59. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema
401 - Unauthorized	Empty

### 10.3.2.3. /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/scale

Table 10.60. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Scale
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.61. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

read scale of the specified StatefulSet

Table 10.62. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
401 - Unauthorized	Empty

## HTTP method

**PATCH**

## Description

partially update scale of the specified StatefulSet

Table 10.63. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 10.64. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 10.65. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace scale of the specified StatefulSet

Table 10.66. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 10.67. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Scale</b> schema	

Table 10.68. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema
401 - Unauthorized	Empty

## 10.3.2.4. /api/v1/namespaces/{namespace}/replicationcontrollers/{name}/scale

Table 10.69. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Scale
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 10.70. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

read scale of the specified ReplicationController

Table 10.71. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
401 - Unauthorized	Empty

## HTTP method

**PATCH**

## Description

partially update scale of the specified ReplicationController

Table 10.72. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed



Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 10.73. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 10.74. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace scale of the specified ReplicationController

**Table 10.75. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 10.76. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Scale</b> schema	

Table 10.77. HTTP responses

HTTP code	Response body
200 - OK	<b>Scale</b> schema
201 - Created	<b>Scale</b> schema
401 - Unauthorized	Empty

## CHAPTER 11. BATCH APIS

### 11.1. BATCH APIS

#### 11.1.1. CronJob [batch/v1]

##### Description

CronJob represents the configuration of a single cron job.

##### Type

**object**

#### 11.1.2. Job [batch/v1]

##### Description

Job represents the configuration of a single job.

##### Type

**object**

### 11.2. CRONJOB [BATCH/V1]

##### Description

CronJob represents the configuration of a single cron job.

##### Type

**object**

#### 11.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	CronJobSpec describes how the job execution will look like and when it will actually run.
<b>status</b>	<b>object</b>	CronJobStatus represents the current state of a cron job.

### 11.2.1.1. .spec

#### Description

CronJobSpec describes how the job execution will look like and when it will actually run.

#### Type

**object**

#### Required

- **schedule**
- **jobTemplate**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>concurrencyPolicy</b>	<b>string</b>	<p>Specifies how to treat concurrent executions of a Job. Valid values are:</p> <ul style="list-style-type: none"> <li>- "Allow" (default): allows CronJobs to run concurrently;</li> <li>- "Forbid": forbids concurrent runs, skipping next run if previous run hasn't finished yet;</li> <li>- "Replace": cancels currently running job and replaces it with a new one</li> </ul> <p>Possible enum values: - <b>"Allow"</b> allows CronJobs to run concurrently. - <b>"Forbid"</b> forbids concurrent runs, skipping next run if previous hasn't finished yet. - <b>"Replace"</b> cancels currently running job and replaces it with a new one.</p>
<b>failedJobsHistoryLimit</b>	<b>integer</b>	The number of failed finished jobs to retain. Value must be non-negative integer. Defaults to 1.
<b>jobTemplate</b>	<b>object</b>	JobTemplateSpec describes the data a Job should have when created from a template
<b>schedule</b>	<b>string</b>	The schedule in Cron format, see <a href="https://en.wikipedia.org/wiki/Cron">https://en.wikipedia.org/wiki/Cron</a> .
<b>startingDeadlineSeconds</b>	<b>integer</b>	Optional deadline in seconds for starting the job if it misses scheduled time for any reason. Missed jobs executions will be counted as failed ones.
<b>successfulJobsHistoryLimit</b>	<b>integer</b>	The number of successful finished jobs to retain. Value must be non-negative integer. Defaults to 3.
<b>suspend</b>	<b>boolean</b>	This flag tells the controller to suspend subsequent executions, it does not apply to already started executions. Defaults to false.

Property	Type	Description
<b>timeZone</b>	<b>string</b>	The time zone name for the given schedule, see <a href="https://en.wikipedia.org/wiki/List_of_tz_database_time_zones">https://en.wikipedia.org/wiki/List_of_tz_database_time_zones</a> . If not specified, this will default to the time zone of the kube-controller-manager process. The set of valid time zone names and the time zone offset is loaded from the system-wide time zone database by the API server during CronJob validation and the controller manager during execution. If no system-wide time zone database can be found a bundled version of the database is used instead. If the time zone name becomes invalid during the lifetime of a CronJob or due to a change in host configuration, the controller will stop creating new new Jobs and will create a system event with the reason UnknownTimeZone. More information can be found in <a href="https://kubernetes.io/docs/concepts/workloads/controllers/cron-jobs/#time-zones">https://kubernetes.io/docs/concepts/workloads/controllers/cron-jobs/#time-zones</a>

### 11.2.1.2. .spec.jobTemplate

#### Description

JobTemplateSpec describes the data a Job should have when created from a template

#### Type

**object**

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata of the jobs created from this template. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	JobSpec describes how the job execution will look like.

### 11.2.1.3. .spec.jobTemplate.spec

#### Description

JobSpec describes how the job execution will look like.

#### Type

**object**

#### Required

- **template**

Property	Type	Description
<b>activeDeadlineSeconds</b>	<b>integer</b>	Specifies the duration in seconds relative to the startTime that the job may be continuously active before the system tries to terminate it; value must be positive integer. If a Job is suspended (at creation or through an update), this timer will effectively be stopped and reset when the Job is resumed again.
<b>backoffLimit</b>	<b>integer</b>	Specifies the number of retries before marking this job failed. Defaults to 6
<b>completionMode</b>	<b>string</b>	<p>completionMode specifies how Pod completions are tracked. It can be <b>NonIndexed</b> (default) or <b>Indexed</b>.</p> <p><b>NonIndexed</b> means that the Job is considered complete when there have been .spec.completions successfully completed Pods. Each Pod completion is homologous to each other.</p> <p><b>Indexed</b> means that the Pods of a Job get an associated completion index from 0 to (.spec.completions - 1), available in the annotation batch.kubernetes.io/job-completion-index. The Job is considered complete when there is one successfully completed Pod for each index. When value is <b>Indexed</b>, .spec.completions must be specified and <b>.spec.parallelism</b> must be less</p>



Property	Type	Description
		<p>than or equal to 10<sup>5</sup>. In addition, The Pod name takes the form <b>\$(job-name)-\$(index)-\$(random-string)</b>, the Pod hostname takes the form <b>\$(job-name)-\$(index)</b>.</p> <p>More completion modes can be added in the future. If the Job controller observes a mode that it doesn't recognize, which is possible during upgrades due to version skew, the controller skips updates for the Job.</p> <p>Possible enum values: -  <b>"Indexed"</b> is a Job completion mode. In this mode, the Pods of a Job get an associated completion index from 0 to (.spec.completions - 1). The Job is considered complete when a Pod completes for each completion index. - <b>"NonIndexed"</b> is a Job completion mode. In this mode, the Job is considered complete when there have been .spec.completions successfully completed Pods. Pod completions are homologous to each other.</p>
<b>completions</b>	<b>integer</b>	<p>Specifies the desired number of successfully finished pods the job should be run with. Setting to null means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job.</p> <p>More info:  <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a></p>

Property	Type	Description
<b>manualSelector</b>	<b>boolean</b>	<p>manualSelector controls generation of pod labels and pod selectors. Leave <b>manualSelector</b> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <b>manualSelector=true</b> in jobs that were created with the old <b>extensions/v1beta1</b> API. More info:  <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector</a></p>
<b>parallelism</b>	<b>integer</b>	<p>Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when <math>((\text{.spec.completions} - \text{.status.successful}) &lt; \text{.spec.parallelism})</math>, i.e. when the work left to do is less than max parallelism. More info:  <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a></p>
<b>podFailurePolicy</b>	<b>object</b>	<p>PodFailurePolicy describes how failed pods influence the backoffLimit.</p>
<b>selector</b>	<b>LabelSelector</b>	<p>A label query over pods that should match the pod count. Normally, the system sets this field for you. More info:  <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors</a></p>

Property	Type	Description
<b>suspend</b>	<b>boolean</b>	suspend specifies whether the Job controller should create Pods or not. If a Job is created with suspend set to true, no Pods are created by the Job controller. If a Job is suspended after creation (i.e. the flag goes from false to true), the Job controller will delete all active Pods associated with this Job. Users must design their workload to gracefully handle this. Suspending a Job will reset the StartTime field of the Job, effectively resetting the ActiveDeadlineSeconds timer too. Defaults to false.
<b>template</b>	<b>PodTemplateSpec</b>	Describes the pod that will be created when executing a job. The only allowed template.spec.restartPolicy values are "Never" or "OnFailure". More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a>
<b>ttlSecondsAfterFinished</b>	<b>integer</b>	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes.

#### 11.2.1.4. .spec.jobTemplate.spec.podFailurePolicy

##### Description

PodFailurePolicy describes how failed pods influence the backoffLimit.

##### Type

**object****Required**

- **rules**

Property	Type	Description
<b>rules</b>	<b>array</b>	A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed.
<b>rules[]</b>	<b>object</b>	PodFailurePolicyRule describes how a pod failure is handled when the requirements are met. One of onExitCodes and onPodConditions, but not both, can be used in each rule.

**11.2.1.5. .spec.jobTemplate.spec.podFailurePolicy.rules****Description**

A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed.

**Type**

**array**

**11.2.1.6. .spec.jobTemplate.spec.podFailurePolicy.rules[]****Description**

PodFailurePolicyRule describes how a pod failure is handled when the requirements are met. One of onExitCodes and onPodConditions, but not both, can be used in each rule.

**Type**

**object**

**Required**

- **action**
- **onPodConditions**

Property	Type	Description
<b>action</b>	<b>string</b>	<p>Specifies the action taken on a pod failure when the requirements are satisfied. Possible values are:</p> <ul style="list-style-type: none"> <li>- FailJob: indicates that the pod's job is marked as Failed and all running pods are terminated.</li> <li>- Ignore: indicates that the counter towards the .backoffLimit is not incremented and a replacement pod is created.</li> <li>- Count: indicates that the pod is handled in the default way - the counter towards the .backoffLimit is incremented.</li> </ul> <p>Additional values are considered to be added in the future. Clients should react to an unknown action by skipping the rule.</p> <p>Possible enum values: - <b>"Count"</b> This is an action which might be taken on a pod failure - the pod failure is handled in the default way - the counter towards .backoffLimit, represented by the job's .status.failed field, is incremented. - <b>"FailJob"</b> This is an action which might be taken on a pod failure - mark the pod's job as Failed and terminate all running pods. - <b>"Ignore"</b> This is an action which might be taken on a pod failure - the counter towards .backoffLimit, represented by the job's .status.failed field, is not incremented and a replacement pod is created.</p>

Property	Type	Description
<b>onExitCodes</b>	<b>object</b>	PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the <code>.state.terminated.exitCode</code> for each app container and init container status, represented by the <code>.status.containerStatuses</code> and <code>.status.initContainerStatuses</code> fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check.
<b>onPodConditions</b>	<b>array</b>	Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed.
<b>onPodConditions[]</b>	<b>object</b>	PodFailurePolicyOnPodConditionPattern describes a pattern for matching an actual pod condition type.

### 11.2.1.7. `.spec.jobTemplate.spec.podFailurePolicy.rules[].onExitCodes`

#### Description

PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the `.state.terminated.exitCode` for each app container and init container status, represented by the `.status.containerStatuses` and `.status.initContainerStatuses` fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check.

#### Type

**object**

#### Required

- **operator**
- **values**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Restricts the check for exit codes to the container with the specified name. When null, the rule applies to all containers. When specified, it should match one the container or initContainer names in the pod template.
<b>operator</b>	<b>string</b>	<p>Represents the relationship between the container exit code(s) and the specified values. Containers completed with success (exit code 0) are excluded from the requirement check. Possible values are:</p> <ul style="list-style-type: none"> <li>- In: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is in the set of specified values.</li> <li>- NotIn: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is not in the set of specified values. Additional values are considered to be added in the future. Clients should react to an unknown operator by assuming the requirement is not satisfied.</li> </ul> <p>Possible enum values: - <b>"In"</b> - <b>"NotIn"</b></p>
<b>values</b>	<b>array (integer)</b>	Specifies the set of values. Each returned container exit code (might be multiple in case of multiple containers) is checked against this set of values with respect to the operator. The list of values must be ordered and must not contain duplicates. Value '0' cannot be used for the In operator. At least one element is required. At most 255 elements are allowed.

### 11.2.1.8. .spec.jobTemplate.spec.podFailurePolicy.rules[].onPodConditions

**Description**

Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed.

**Type**

**array**

**11.2.1.9. .spec.jobTemplate.spec.podFailurePolicy.rules[].onPodConditions[]****Description**

PodFailurePolicyOnPodConditionsPattern describes a pattern for matching an actual pod condition type.

**Type**

**object**

**Required**

- **type**
- **status**

Property	Type	Description
<b>status</b>	<b>string</b>	Specifies the required Pod condition status. To match a pod condition it is required that the specified status equals the pod condition status. Defaults to True.
<b>type</b>	<b>string</b>	Specifies the required Pod condition type. To match a pod condition it is required that specified type equals the pod condition type.

**11.2.1.10. .status****Description**

CronJobStatus represents the current state of a cron job.

**Type**

**object**

Property	Type	Description
<b>active</b>	<b>array (ObjectReference)</b>	A list of pointers to currently running jobs.



Property	Type	Description
<b>lastScheduleTime</b>	<b>Time</b>	Information when was the last time the job was successfully scheduled.
<b>lastSuccessfulTime</b>	<b>Time</b>	Information when was the last time the job successfully completed.

### 11.2.2. API endpoints

The following API endpoints are available:

- **/apis/batch/v1/cronjobs**
  - **GET**: list or watch objects of kind CronJob
- **/apis/batch/v1/watch/cronjobs**
  - **GET**: watch individual changes to a list of CronJob. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/batch/v1/namespaces/{namespace}/cronjobs**
  - **DELETE**: delete collection of CronJob
  - **GET**: list or watch objects of kind CronJob
  - **POST**: create a CronJob
- **/apis/batch/v1/watch/namespaces/{namespace}/cronjobs**
  - **GET**: watch individual changes to a list of CronJob. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/batch/v1/namespaces/{namespace}/cronjobs/{name}**
  - **DELETE**: delete a CronJob
  - **GET**: read the specified CronJob
  - **PATCH**: partially update the specified CronJob
  - **PUT**: replace the specified CronJob
- **/apis/batch/v1/watch/namespaces/{namespace}/cronjobs/{name}**
  - **GET**: watch changes to an object of kind CronJob. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status**

- **GET**: read status of the specified CronJob
- **PATCH**: partially update status of the specified CronJob
- **PUT**: replace status of the specified CronJob

### 11.2.2.1. /apis/batch/v1/cronjobs

Table 11.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind CronJob

Table 11.2. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CronJobList</a> schema
401 - Unauthorized	Empty

### 11.2.2.2. /apis/batch/v1/watch/cronjobs

Table 11.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CronJob. deprecated: use the 'watch' parameter with a list operation instead.

Table 11.4. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 11.2.2.3. /apis/batch/v1/namespaces/{namespace}/cronjobs

Table 11.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of CronJob

Table 11.7. Query parameters

Parameter	Type	Description
-----------	------	-------------



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 11.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 11.9. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind CronJob

Table 11.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 11.11. HTTP responses

HTTP code	Response body
200 - OK	<b>CronJobList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a CronJob

**Table 11.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 11.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CronJob</b> schema	

**Table 11.14. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>CronJob</b> schema



HTTP code	Reponse body
201 - Created	<b>CronJob</b> schema
202 - Accepted	<b>CronJob</b> schema
401 - Unauthorized	Empty

#### 11.2.2.4. /apis/batch/v1/watch/namespaces/{namespace}/cronjobs

Table 11.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CronJob. deprecated: use the 'watch' parameter with a list operation instead.

Table 11.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 11.2.2.5. /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}

Table 11.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CronJob
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a CronJob

Table 11.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 11.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 11.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified CronJob

Table 11.23. HTTP responses

HTTP code	Response body
200 - OK	<b>CronJob</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified CronJob

**Table 11.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 11.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 11.26. HTTP responses

HTTP code	Reponse body
200 - OK	<b>CronJob</b> schema
201 - Created	<b>CronJob</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified CronJob

Table 11.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 11.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CronJob</b> schema	

Table 11.29. HTTP responses

HTTP code	Response body
200 - OK	<b>CronJob</b> schema
201 - Created	<b>CronJob</b> schema
401 - Unauthorized	Empty

### 11.2.2.6. /apis/batch/v1/watch/namespaces/{namespace}/cronjobs/{name}

Table 11.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CronJob
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind CronJob. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 11.32. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 11.2.2.7. /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status

Table 11.33. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CronJob
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.34. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified CronJob

Table 11.35. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">CronJob</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified CronJob

Table 11.36. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 11.37. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 11.38. HTTP responses

HTTP code	Response body
200 - OK	<b>CronJob</b> schema
201 - Created	<b>CronJob</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified CronJob

**Table 11.39. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 11.40. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CronJob</b> schema	

Table 11.41. HTTP responses

HTTP code	Response body
200 - OK	<b>CronJob</b> schema
201 - Created	<b>CronJob</b> schema
401 - Unauthorized	Empty

## 11.3. JOB [BATCH/V1]

### Description

Job represents the configuration of a single job.

### Type

**object**

### 11.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>



Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	JobSpec describes how the job execution will look like.
<b>status</b>	<b>object</b>	JobStatus represents the current state of a Job.

### 11.3.1.1. .spec

#### Description

JobSpec describes how the job execution will look like.

#### Type

**object**

#### Required

- **template**

Property	Type	Description
<b>activeDeadlineSeconds</b>	<b>integer</b>	Specifies the duration in seconds relative to the startTime that the job may be continuously active before the system tries to terminate it; value must be positive integer. If a Job is suspended (at creation or through an update), this timer will effectively be stopped and reset when the Job is resumed again.
<b>backoffLimit</b>	<b>integer</b>	Specifies the number of retries before marking this job failed. Defaults to 6
<b>completionMode</b>	<b>string</b>	completionMode specifies how Pod completions are tracked. It can be <b>NonIndexed</b> (default) or <b>Indexed</b> .  <b>NonIndexed</b> means that the

Property	Type	Description
		<p>Job is considered complete when there have been <code>.spec.completions</code> successfully completed Pods. Each Pod completion is homologous to each other.</p> <p><b>Indexed</b> means that the Pods of a Job get an associated completion index from 0 to <math>(.spec.completions - 1)</math>, available in the annotation <code>batch.kubernetes.io/job-completion-index</code>. The Job is considered complete when there is one successfully completed Pod for each index. When value is <b>Indexed</b>, <code>.spec.completions</code> must be specified and <code>.spec.parallelism</code> must be less than or equal to <math>10^5</math>. In addition, The Pod name takes the form <b><code>\$(job-name)-\$(index)-\$(random-string)</code></b>, the Pod hostname takes the form <b><code>\$(job-name)-\$(index)</code></b>.</p> <p>More completion modes can be added in the future. If the Job controller observes a mode that it doesn't recognize, which is possible during upgrades due to version skew, the controller skips updates for the Job.</p> <p>Possible enum values: -  <b>"Indexed"</b> is a Job completion mode. In this mode, the Pods of a Job get an associated completion index from 0 to <math>(.spec.completions - 1)</math>. The Job is considered complete when a Pod completes for each completion index. - <b>"NonIndexed"</b> is a Job completion mode. In this mode, the Job is considered complete when there have been <code>.spec.completions</code> successfully completed Pods. Pod completions are homologous to each other.</p>

Property	Type	Description
<b>completions</b>	<b>integer</b>	<p>Specifies the desired number of successfully finished pods the job should be run with. Setting to null means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job.</p> <p>More info:  <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a></p>
<b>manualSelector</b>	<b>boolean</b>	<p><code>manualSelector</code> controls generation of pod labels and pod selectors. Leave <b><code>manualSelector</code></b> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <b><code>manualSelector=true</code></b> in jobs that were created with the old <b><code>extensions/v1beta1</code></b> API. More info:  <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector</a></p>

Property	Type	Description
<b>parallelism</b>	<b>integer</b>	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$ , i.e. when the work left to do is less than max parallelism. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a>
<b>podFailurePolicy</b>	<b>object</b>	PodFailurePolicy describes how failed pods influence the backoffLimit.
<b>selector</b>	<b>LabelSelector</b>	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors</a>
<b>suspend</b>	<b>boolean</b>	suspend specifies whether the Job controller should create Pods or not. If a Job is created with suspend set to true, no Pods are created by the Job controller. If a Job is suspended after creation (i.e. the flag goes from false to true), the Job controller will delete all active Pods associated with this Job. Users must design their workload to gracefully handle this. Suspending a Job will reset the StartTime field of the Job, effectively resetting the ActiveDeadlineSeconds timer too. Defaults to false.

Property	Type	Description
<b>template</b>	<b>PodTemplateSpec</b>	Describes the pod that will be created when executing a job. The only allowed <code>template.spec.restartPolicy</code> values are "Never" or "OnFailure". More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a>
<b>ttlSecondsAfterFinished</b>	<b>integer</b>	<code>ttlSecondsAfterFinished</code> limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, <code>ttlSecondsAfterFinished</code> after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes.

### 11.3.1.2. .spec.podFailurePolicy

#### Description

`PodFailurePolicy` describes how failed pods influence the `backoffLimit`.

#### Type

**object**

#### Required

- **rules**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>rules</b>	<b>array</b>	A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed.
<b>rules[]</b>	<b>object</b>	PodFailurePolicyRule describes how a pod failure is handled when the requirements are met. One of onExitCodes and onPodConditions, but not both, can be used in each rule.

### 11.3.1.3. .spec.podFailurePolicy.rules

#### Description

A list of pod failure policy rules. The rules are evaluated in order. Once a rule matches a Pod failure, the remaining of the rules are ignored. When no rule matches the Pod failure, the default handling applies - the counter of pod failures is incremented and it is checked against the backoffLimit. At most 20 elements are allowed.

#### Type

**array**

### 11.3.1.4. .spec.podFailurePolicy.rules[]

#### Description

PodFailurePolicyRule describes how a pod failure is handled when the requirements are met. One of onExitCodes and onPodConditions, but not both, can be used in each rule.

#### Type

**object**

#### Required

- **action**
- **onPodConditions**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>action</b>	<b>string</b>	<p>Specifies the action taken on a pod failure when the requirements are satisfied. Possible values are:</p> <ul style="list-style-type: none"> <li>- FailJob: indicates that the pod's job is marked as Failed and all running pods are terminated.</li> <li>- Ignore: indicates that the counter towards the .backoffLimit is not incremented and a replacement pod is created.</li> <li>- Count: indicates that the pod is handled in the default way - the counter towards the .backoffLimit is incremented.</li> </ul> <p>Additional values are considered to be added in the future. Clients should react to an unknown action by skipping the rule.</p> <p>Possible enum values: - <b>"Count"</b> This is an action which might be taken on a pod failure - the pod failure is handled in the default way - the counter towards .backoffLimit, represented by the job's .status.failed field, is incremented. - <b>"FailJob"</b> This is an action which might be taken on a pod failure - mark the pod's job as Failed and terminate all running pods. - <b>"Ignore"</b> This is an action which might be taken on a pod failure - the counter towards .backoffLimit, represented by the job's .status.failed field, is not incremented and a replacement pod is created.</p>

Property	Type	Description
<b>onExitCodes</b>	<b>object</b>	PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the <code>.state.terminated.exitCode</code> for each app container and init container status, represented by the <code>.status.containerStatuses</code> and <code>.status.initContainerStatuses</code> fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check.
<b>onPodConditions</b>	<b>array</b>	Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed.
<b>onPodConditions[]</b>	<b>object</b>	PodFailurePolicyOnPodConditionPattern describes a pattern for matching an actual pod condition type.

### 11.3.1.5. `.spec.podFailurePolicy.rules[].onExitCodes`

#### Description

PodFailurePolicyOnExitCodesRequirement describes the requirement for handling a failed pod based on its container exit codes. In particular, it lookups the `.state.terminated.exitCode` for each app container and init container status, represented by the `.status.containerStatuses` and `.status.initContainerStatuses` fields in the Pod status, respectively. Containers completed with success (exit code 0) are excluded from the requirement check.

#### Type

**object**

#### Required

- **operator**
- **values**



Property	Type	Description
<b>containerName</b>	<b>string</b>	Restricts the check for exit codes to the container with the specified name. When null, the rule applies to all containers. When specified, it should match one the container or initContainer names in the pod template.
<b>operator</b>	<b>string</b>	<p>Represents the relationship between the container exit code(s) and the specified values. Containers completed with success (exit code 0) are excluded from the requirement check. Possible values are:</p> <ul style="list-style-type: none"> <li>- In: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is in the set of specified values.</li> <li>- NotIn: the requirement is satisfied if at least one container exit code (might be multiple if there are multiple containers not restricted by the 'containerName' field) is not in the set of specified values. Additional values are considered to be added in the future. Clients should react to an unknown operator by assuming the requirement is not satisfied.</li> </ul> <p>Possible enum values: - <b>"In"</b> - <b>"NotIn"</b></p>
<b>values</b>	<b>array (integer)</b>	Specifies the set of values. Each returned container exit code (might be multiple in case of multiple containers) is checked against this set of values with respect to the operator. The list of values must be ordered and must not contain duplicates. Value '0' cannot be used for the In operator. At least one element is required. At most 255 elements are allowed.

### 11.3.1.6. .spec.podFailurePolicy.rules[].onPodConditions

**Description**

Represents the requirement on the pod conditions. The requirement is represented as a list of pod condition patterns. The requirement is satisfied if at least one pattern matches an actual pod condition. At most 20 elements are allowed.

**Type**

**array**

**11.3.1.7. .spec.podFailurePolicy.rules[].onPodConditions[]****Description**

PodFailurePolicyOnPodConditionsPattern describes a pattern for matching an actual pod condition type.

**Type**

**object**

**Required**

- **type**
- **status**

Property	Type	Description
<b>status</b>	<b>string</b>	Specifies the required Pod condition status. To match a pod condition it is required that the specified status equals the pod condition status. Defaults to True.
<b>type</b>	<b>string</b>	Specifies the required Pod condition type. To match a pod condition it is required that specified type equals the pod condition type.

**11.3.1.8. .status****Description**

JobStatus represents the current state of a Job.

**Type**

**object**

Property	Type	Description
<b>active</b>	<b>integer</b>	The number of pending and running pods.

Property	Type	Description
<b>completedIndexes</b>	<b>string</b>	completedIndexes holds the completed indexes when <code>.spec.completionMode = "Indexed"</code> in a text format. The indexes are represented as decimal integers separated by commas. The numbers are listed in increasing order. Three or more consecutive numbers are compressed and represented by the first and last element of the series, separated by a hyphen. For example, if the completed indexes are 1, 3, 4, 5 and 7, they are represented as "1,3-5,7".
<b>completionTime</b>	<b>Time</b>	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC. The completion time is only set when the job finishes successfully.
<b>conditions</b>	<b>array</b>	The latest available observations of an object's current state. When a Job fails, one of the conditions will have type "Failed" and status true. When a Job is suspended, one of the conditions will have type "Suspended" and status true; when the Job is resumed, the status of this condition will become false. When a Job is completed, one of the conditions will have type "Complete" and status true. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/">https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/</a>
<b>conditions[]</b>	<b>object</b>	JobCondition describes current state of a job.
<b>failed</b>	<b>integer</b>	The number of pods which reached phase Failed.

Property	Type	Description
<b>ready</b>	<b>integer</b>	The number of pods which have a Ready condition.  This field is beta-level. The job controller populates the field when the feature gate JobReadyPods is enabled (enabled by default).
<b>startTime</b>	<b>Time</b>	Represents time when the job controller started processing a job. When a Job is created in the suspended state, this field is not set until the first time it is resumed. This field is reset every time a Job is resumed from suspension. It is represented in RFC3339 form and is in UTC.
<b>succeeded</b>	<b>integer</b>	The number of pods which reached phase Succeeded.
<b>uncountedTerminatedPods</b>	<b>object</b>	UncountedTerminatedPods holds UIDs of Pods that have terminated but haven't been accounted in Job status counters.

### 11.3.1.9. .status.conditions

#### Description

The latest available observations of an object's current state. When a Job fails, one of the conditions will have type "Failed" and status true. When a Job is suspended, one of the conditions will have type "Suspended" and status true; when the Job is resumed, the status of this condition will become false. When a Job is completed, one of the conditions will have type "Complete" and status true. More info: <https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/>

#### Type

**array**

### 11.3.1.10. .status.conditions[]

#### Description

JobCondition describes current state of a job.

#### Type

**object**

#### Required

- **type**

- **status**

Property	Type	Description
<b>lastProbeTime</b>	<b>Time</b>	Last time the condition was checked.
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transit from one status to another.
<b>message</b>	<b>string</b>	Human readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	(brief) reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of job condition, Complete or Failed.

### 11.3.1.11. .status.uncountedTerminatedPods

#### Description

UncountedTerminatedPods holds UIDs of Pods that have terminated but haven't been accounted in Job status counters.

#### Type

**object**

Property	Type	Description
<b>failed</b>	<b>array (string)</b>	failed holds UIDs of failed Pods.
<b>succeeded</b>	<b>array (string)</b>	succeeded holds UIDs of succeeded Pods.

### 11.3.2. API endpoints

The following API endpoints are available:

- **/apis/batch/v1/jobs**
  - **GET**: list or watch objects of kind Job
- **/apis/batch/v1/watch/jobs**

- **GET**: watch individual changes to a list of Job. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/batch/v1/namespaces/{namespace}/jobs**
  - **DELETE**: delete collection of Job
  - **GET**: list or watch objects of kind Job
  - **POST**: create a Job
- **/apis/batch/v1/watch/namespaces/{namespace}/jobs**
  - **GET**: watch individual changes to a list of Job. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/batch/v1/namespaces/{namespace}/jobs/{name}**
  - **DELETE**: delete a Job
  - **GET**: read the specified Job
  - **PATCH**: partially update the specified Job
  - **PUT**: replace the specified Job
- **/apis/batch/v1/watch/namespaces/{namespace}/jobs/{name}**
  - **GET**: watch changes to an object of kind Job. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/batch/v1/namespaces/{namespace}/jobs/{name}/status**
  - **GET**: read status of the specified Job
  - **PATCH**: partially update status of the specified Job
  - **PUT**: replace status of the specified Job

### 11.3.2.1. /apis/batch/v1/jobs

Table 11.42. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Job

**Table 11.43. HTTP responses**

HTTP code	Response body
200 - OK	<b>JobList</b> schema
401 - Unauthorized	Empty

### 11.3.2.2. /apis/batch/v1/watch/jobs

Table 11.44. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Job. deprecated: use the 'watch' parameter with a list operation instead.

Table 11.45. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 11.3.2.3. /apis/batch/v1/namespaces/{namespace}/jobs

Table 11.46. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.47. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Job

Table 11.48. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>



Table 11.49. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 11.50. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Job

Table 11.51. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 11.52. HTTP responses

HTTP code	Response body
200 - OK	<b>JobList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a Job

**Table 11.53. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 11.54. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Job</b> schema	

**Table 11.55. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Job</b> schema

HTTP code	Response body
201 - Created	<b>Job</b> schema
202 - Accepted	<b>Job</b> schema
401 - Unauthorized	Empty

#### 11.3.2.4. /apis/batch/v1/watch/namespaces/{namespace}/jobs

Table 11.56. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.57. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Job. deprecated: use the 'watch' parameter with a list operation instead.

Table 11.58. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 11.3.2.5. /apis/batch/v1/namespaces/{namespace}/jobs/{name}

Table 11.59. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Job
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.60. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a Job

Table 11.61. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 11.62. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 11.63. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Job

Table 11.64. HTTP responses

HTTP code	Response body
200 - OK	<b>Job</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Job

**Table 11.65. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 11.66. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 11.67. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Job</b> schema
201 - Created	<b>Job</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Job

Table 11.68. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 11.69. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Job</b> schema	

Table 11.70. HTTP responses

HTTP code	Response body
200 - OK	<b>Job</b> schema
201 - Created	<b>Job</b> schema
401 - Unauthorized	Empty

### 11.3.2.6. /apis/batch/v1/watch/namespaces/{namespace}/jobs/{name}

Table 11.71. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Job
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.72. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Job. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 11.73. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 11.3.2.7. /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status

Table 11.74. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Job
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 11.75. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Job

Table 11.76. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">Job</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Job

Table 11.77. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 11.78. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 11.79. HTTP responses

HTTP code	Response body
200 - OK	<b>Job</b> schema
201 - Created	<b>Job</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Job

**Table 11.80. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 11.81. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Job</b> schema	

Table 11.82. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Job</b> schema
201 - Created	<b>Job</b> schema
401 - Unauthorized	Empty

## CHAPTER 12. CERTIFICATES APIS

### 12.1. CERTIFICATES APIS

#### 12.1.1. CertificateSigningRequest [certificates.k8s.io/v1]

##### Description

CertificateSigningRequest objects provide a mechanism to obtain x509 certificates by submitting a certificate signing request, and having it asynchronously approved and issued.

Kubelets use this API to obtain: 1. client certificates to authenticate to kube-apiserver (with the "kubernetes.io/kube-apiserver-client-kubelet" signerName). 2. serving certificates for TLS endpoints kube-apiserver can connect to securely (with the "kubernetes.io/kubelet-serving" signerName).

This API can be used to request client certificates to authenticate to kube-apiserver (with the "kubernetes.io/kube-apiserver-client" signerName), or to obtain certificates from custom non-Kubernetes signers.

##### Type

**object**

### 12.2. CERTIFICATESIGNINGREQUEST [CERTIFICATES.K8S.IO/V1]

##### Description

CertificateSigningRequest objects provide a mechanism to obtain x509 certificates by submitting a certificate signing request, and having it asynchronously approved and issued.

Kubelets use this API to obtain: 1. client certificates to authenticate to kube-apiserver (with the "kubernetes.io/kube-apiserver-client-kubelet" signerName). 2. serving certificates for TLS endpoints kube-apiserver can connect to securely (with the "kubernetes.io/kubelet-serving" signerName).

This API can be used to request client certificates to authenticate to kube-apiserver (with the "kubernetes.io/kube-apiserver-client" signerName), or to obtain certificates from custom non-Kubernetes signers.

##### Type

**object**

##### Required

- **spec**

#### 12.2.1. Specification

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	
<b>spec</b>	<b>object</b>	CertificateSigningRequestSpec contains the certificate request.
<b>status</b>	<b>object</b>	CertificateSigningRequestStatus contains conditions used to indicate approved/denied/failed status of the request, and the issued certificate.

### 12.2.1.1. .spec

#### Description

CertificateSigningRequestSpec contains the certificate request.

#### Type

**object**

#### Required

- **request**
- **signerName**

Property	Type	Description
<b>expirationSeconds</b>	<b>integer</b>	<p>expirationSeconds is the requested duration of validity of the issued certificate. The certificate signer may issue a certificate with a different validity duration so a client must check the delta between the notBefore and and notAfter fields in the issued certificate to determine the actual duration.</p> <p>The v1.22+ in-tree implementations of the well-known Kubernetes signers will honor this field as long as the requested duration is not greater than the maximum duration they will honor per the --cluster-signing-duration CLI flag to the Kubernetes controller manager.</p> <p>Certificate signers may not honor this field for various reasons:</p> <ol style="list-style-type: none"> <li>1. Old signer that is unaware of the field (such as the in-tree implementations prior to v1.22)</li> <li>2. Signer whose configured maximum is shorter than the requested duration</li> <li>3. Signer whose configured minimum is longer than the requested duration</li> </ol> <p>The minimum valid value for expirationSeconds is 600, i.e. 10 minutes.</p>
<b>extra</b>	<b>object</b>	<p>extra contains extra attributes of the user that created the CertificateSigningRequest. Populated by the API server on creation and immutable.</p>
<b>extra{}</b>	<b>array (string)</b>	
<b>groups</b>	<b>array (string)</b>	<p>groups contains group membership of the user that created the CertificateSigningRequest. Populated by the API server on creation and immutable.</p>



Property	Type	Description
<b>request</b>	<b>string</b>	request contains an x509 certificate signing request encoded in a "CERTIFICATE REQUEST" PEM block. When serialized as JSON or YAML, the data is additionally base64-encoded.
<b>signerName</b>	<b>string</b>	<p>signerName indicates the requested signer, and is a qualified name.</p> <p>List/watch requests for CertificateSigningRequests can filter on this field using a "spec.signerName=NAME" fieldSelector.</p> <p>Well-known Kubernetes signers are: 1. "kubernetes.io/kube-apiserver-client": issues client certificates that can be used to authenticate to kube-apiserver. Requests for this signer are never auto-approved by kube-controller-manager, can be issued by the "csrsigning" controller in kube-controller-manager. 2. "kubernetes.io/kube-apiserver-client-kubelet": issues client certificates that kubelets use to authenticate to kube-apiserver. Requests for this signer can be auto-approved by the "csrapproving" controller in kube-controller-manager, and can be issued by the "csrsigning" controller in kube-controller-manager. 3. "kubernetes.io/kubelet-serving" issues serving certificates that kubelets use to serve TLS endpoints, which kube-apiserver can connect to securely. Requests for this signer are never auto-approved by kube-controller-manager, and can be issued by the "csrsigning" controller in kube-controller-manager.</p> <p>More details are available at</p>

Property	Type	Description
		<p><a href="https://k8s.io/docs/reference/access-authorization/certificate-signing-requests/#kubernetes-signers">https://k8s.io/docs/reference/access-authorization/certificate-signing-requests/#kubernetes-signers</a></p> <p>Custom signerNames can also be specified. The signer defines:</p> <ol style="list-style-type: none"> <li>Trust distribution: how trust (CA bundles) are distributed.</li> <li>Permitted subjects: and behavior when a disallowed subject is requested.</li> <li>Required, permitted, or forbidden x509 extensions in the request (including whether subjectAltNames are allowed, which types, restrictions on allowed values) and behavior when a disallowed extension is requested.</li> <li>Required, permitted, or forbidden key usages / extended key usages.</li> <li>Expiration/certificate lifetime: whether it is fixed by the signer, configurable by the admin.</li> <li>Whether or not requests for CA certificates are allowed.</li> </ol>
<b>uid</b>	<b>string</b>	<p>uid contains the uid of the user that created the CertificateSigningRequest. Populated by the API server on creation and immutable.</p>

Property	Type	Description
<b>usages</b>	<b>array (string)</b>	<p>usages specifies a set of key usages requested in the issued certificate.</p> <p>Requests for TLS client certificates typically request: "digital signature", "key encipherment", "client auth".</p> <p>Requests for TLS serving certificates typically request: "key encipherment", "digital signature", "server auth".</p> <p>Valid values are: "signing", "digital signature", "content commitment", "key encipherment", "key agreement", "data encipherment", "cert sign", "crl sign", "encipher only", "decipher only", "any", "server auth", "client auth", "code signing", "email protection", "s/mime", "ipsec end system", "ipsec tunnel", "ipsec user", "timestamping", "ocsp signing", "microsoft sgc", "netscape sgc"</p>
<b>username</b>	<b>string</b>	<p>username contains the name of the user that created the CertificateSigningRequest. Populated by the API server on creation and immutable.</p>

### 12.2.1.2. .spec.extra

#### Description

extra contains extra attributes of the user that created the CertificateSigningRequest. Populated by the API server on creation and immutable.

#### Type

**object**

### 12.2.1.3. .status

#### Description

CertificateSigningRequestStatus contains conditions used to indicate approved/denied/failed status of the request, and the issued certificate.

#### Type

**object**

Property	Type	Description
<b>certificate</b>	<b>string</b>	<p>certificate is populated with an issued certificate by the signer after an Approved condition is present. This field is set via the /status subresource. Once populated, this field is immutable.</p> <p>If the certificate signing request is denied, a condition of type "Denied" is added and this field remains empty. If the signer cannot issue the certificate, a condition of type "Failed" is added and this field remains empty.</p> <p>Validation requirements: 1. certificate must contain one or more PEM blocks. 2. All PEM blocks must have the "CERTIFICATE" label, contain no headers, and the encoded data must be a BER-encoded ASN.1 Certificate structure as described in section 4 of RFC5280. 3. Non-PEM content may appear before or after the "CERTIFICATE" PEM blocks and is unvalidated, to allow for explanatory text as described in section 5.2 of RFC7468.</p> <p>If more than one PEM block is present, and the definition of the requested spec.signerName does not indicate otherwise, the first block is the issued certificate, and subsequent blocks should be treated as intermediate certificates and presented in TLS handshakes.</p> <p>The certificate is encoded in PEM format.</p> <p>When serialized as JSON or YAML, the data is additionally base64-encoded, so it consists of:</p> <pre>base64( -----BEGIN CERTIFICATE----- ... -----END CERTIFICATE----- )</pre>

Property	Type	Description
<b>conditions</b>	<b>array</b>	conditions applied to the request. Known conditions are "Approved", "Denied", and "Failed".
<b>conditions[]</b>	<b>object</b>	CertificateSigningRequestCondition describes a condition of a CertificateSigningRequest object

#### 12.2.1.4. .status.conditions

##### Description

conditions applied to the request. Known conditions are "Approved", "Denied", and "Failed".

##### Type

**array**

#### 12.2.1.5. .status.conditions[]

##### Description

CertificateSigningRequestCondition describes a condition of a CertificateSigningRequest object

##### Type

**object**

##### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	lastTransitionTime is the time the condition last transitioned from one status to another. If unset, when a new condition type is added or an existing condition's status is changed, the server defaults this to the current time.
<b>lastUpdateTime</b>	<b>Time</b>	lastUpdateTime is the time of the last update to this condition
<b>message</b>	<b>string</b>	message contains a human readable message with details about the request state

Property	Type	Description
<b>reason</b>	<b>string</b>	reason indicates a brief reason for the request state
<b>status</b>	<b>string</b>	status of the condition, one of True, False, Unknown. Approved, Denied, and Failed conditions may not be "False" or "Unknown".
<b>type</b>	<b>string</b>	<p>type of the condition. Known conditions are "Approved", "Denied", and "Failed".</p> <p>An "Approved" condition is added via the /approval subresource, indicating the request was approved and should be issued by the signer.</p> <p>A "Denied" condition is added via the /approval subresource, indicating the request was denied and should not be issued by the signer.</p> <p>A "Failed" condition is added via the /status subresource, indicating the signer failed to issue the certificate.</p> <p>Approved and Denied conditions are mutually exclusive. Approved, Denied, and Failed conditions cannot be removed once added.</p> <p>Only one condition of a given type is allowed.</p>

### 12.2.2. API endpoints

The following API endpoints are available:

- **/apis/certificates.k8s.io/v1/certificatesigningrequests**
  - **DELETE**: delete collection of CertificateSigningRequest
  - **GET**: list or watch objects of kind CertificateSigningRequest
  - **POST**: create a CertificateSigningRequest
- **/apis/certificates.k8s.io/v1/watch/certificatesigningrequests**
  - **GET**: watch individual changes to a list of CertificateSigningRequest. deprecated: use the 'watch' parameter with a list operation instead

`watch` parameter with a list operation instead.

- **/apis/certificates.k8s.io/v1/certificatesigningrequests/{name}**
  - **DELETE:** delete a CertificateSigningRequest
  - **GET:** read the specified CertificateSigningRequest
  - **PATCH:** partially update the specified CertificateSigningRequest
  - **PUT:** replace the specified CertificateSigningRequest
- **/apis/certificates.k8s.io/v1/watch/certificatesigningrequests/{name}**
  - **GET:** watch changes to an object of kind CertificateSigningRequest. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/certificates.k8s.io/v1/certificatesigningrequests/{name}/status**
  - **GET:** read status of the specified CertificateSigningRequest
  - **PATCH:** partially update status of the specified CertificateSigningRequest
  - **PUT:** replace status of the specified CertificateSigningRequest
- **/apis/certificates.k8s.io/v1/certificatesigningrequests/{name}/approval**
  - **GET:** read approval of the specified CertificateSigningRequest
  - **PATCH:** partially update approval of the specified CertificateSigningRequest
  - **PUT:** replace approval of the specified CertificateSigningRequest

### 12.2.2.1. /apis/certificates.k8s.io/v1/certificatesigningrequests

Table 12.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**DELETE**

Description

delete collection of CertificateSigningRequest

Table 12.2. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 12.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 12.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind CertificateSigningRequest

Table 12.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 12.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CertificateSigningRequestList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a CertificateSigningRequest

**Table 12.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 12.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CertificateSigningRequest</b> schema	

**Table 12.9. HTTP responses**

HTTP code	Response body
200 - OK	<b>CertificateSigningRequest</b> schema

HTTP code	Response body
201 - Created	<a href="#">CertificateSigningRequest</a> schema
202 - Accepted	<a href="#">CertificateSigningRequest</a> schema
401 - Unauthorized	Empty

### 12.2.2.2. /apis/certificates.k8s.io/v1/watch/certificateSigningRequests

Table 12.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.



Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CertificateSigningRequest. deprecated: use the 'watch' parameter with a list operation instead.

Table 12.11. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 12.2.2.3. /apis/certificates.k8s.io/v1/certificatesigningrequests/{name}

Table 12.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CertificateSigningRequest

Table 12.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a CertificateSigningRequest

Table 12.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 12.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 12.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified CertificateSigningRequest

Table 12.17. HTTP responses

HTTP code	Response body
200 - OK	<b>CertificateSigningRequest</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified CertificateSigningRequest

**Table 12.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 12.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 12.20. HTTP responses

HTTP code	Response body
200 - OK	<b>CertificateSigningRequest</b> schema
201 - Created	<b>CertificateSigningRequest</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified CertificateSigningRequest

Table 12.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 12.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CertificateSigningRequest</b> schema	

Table 12.23. HTTP responses

HTTP code	Response body
200 - OK	<b>CertificateSigningRequest</b> schema
201 - Created	<b>CertificateSigningRequest</b> schema
401 - Unauthorized	Empty

#### 12.2.2.4. /apis/certificates.k8s.io/v1/watch/certificatesigningrequests/{name}

Table 12.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CertificateSigningRequest

Table 12.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind CertificateSigningRequest. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 12.26. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 12.2.2.5. /apis/certificates.k8s.io/v1/certificatesigningrequests/{name}/status

Table 12.27. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CertificateSigningRequest

Table 12.28. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **GET**

#### Description

read status of the specified CertificateSigningRequest

Table 12.29. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">CertificateSigningRequest</a> schema
401 - Unauthorized	Empty

#### HTTP method

#### **PATCH**

#### Description

partially update status of the specified CertificateSigningRequest

Table 12.30. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 12.31. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 12.32. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CertificateSigningRequest</a> schema
201 - Created	<a href="#">CertificateSigningRequest</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified CertificateSigningRequest

**Table 12.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 12.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<a href="#">CertificateSigningRequest</a> schema	

Table 12.35. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">CertificateSigningRequest</a> schema
201 - Created	<a href="#">CertificateSigningRequest</a> schema
401 - Unauthorized	Empty

### 12.2.2.6. /apis/certificates.k8s.io/v1/certificatesigningrequests/{name}/approval

Table 12.36. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CertificateSigningRequest

Table 12.37. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read approval of the specified CertificateSigningRequest

Table 12.38. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">CertificateSigningRequest</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update approval of the specified CertificateSigningRequest

Table 12.39. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 12.40. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 12.41. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CertificateSigningRequest</a> schema
201 - Created	<a href="#">CertificateSigningRequest</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace approval of the specified [CertificateSigningRequest](#)

Table 12.42. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized <code>dryRun</code> directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	<code>fieldManager</code> is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	<code>fieldValidation</code> instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a <code>BadRequest</code> error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.



Table 12.43. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CertificateSigningRequest</b> schema	

Table 12.44. HTTP responses

HTTP code	Response body
200 - OK	<b>CertificateSigningRequest</b> schema
201 - Created	<b>CertificateSigningRequest</b> schema
401 - Unauthorized	Empty

## CHAPTER 13. COORDINATION APIS

### 13.1. COORDINATION APIS

#### 13.1.1. Lease [coordination.k8s.io/v1]

##### Description

Lease defines a lease concept.

##### Type

**object**

### 13.2. LEASE [COORDINATION.K8S.IO/V1]

##### Description

Lease defines a lease concept.

##### Type

**object**

#### 13.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	LeaseSpec is a specification of a Lease.

### 13.2.1.1. .spec

#### Description

LeaseSpec is a specification of a Lease.

#### Type

**object**

Property	Type	Description
<b>acquireTime</b>	<b>MicroTime</b>	acquireTime is a time when the current lease was acquired.
<b>holderIdentity</b>	<b>string</b>	holderIdentity contains the identity of the holder of a current lease.
<b>leaseDurationSeconds</b>	<b>integer</b>	leaseDurationSeconds is a duration that candidates for a lease need to wait to force acquire it. This is measure against time of last observed renewTime.
<b>leaseTransitions</b>	<b>integer</b>	leaseTransitions is the number of transitions of a lease between holders.
<b>renewTime</b>	<b>MicroTime</b>	renewTime is a time when the current holder of a lease has last updated the lease.

### 13.2.2. API endpoints

The following API endpoints are available:

- **/apis/coordination.k8s.io/v1/leases**
  - **GET**: list or watch objects of kind Lease
- **/apis/coordination.k8s.io/v1/watch/leases**

- **GET**: watch individual changes to a list of Lease. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/coordination.k8s.io/v1/namespaces/{namespace}/leases**
  - **DELETE**: delete collection of Lease
  - **GET**: list or watch objects of kind Lease
  - **POST**: create a Lease
- **/apis/coordination.k8s.io/v1/watch/namespaces/{namespace}/leases**
  - **GET**: watch individual changes to a list of Lease. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/coordination.k8s.io/v1/namespaces/{namespace}/leases/{name}**
  - **DELETE**: delete a Lease
  - **GET**: read the specified Lease
  - **PATCH**: partially update the specified Lease
  - **PUT**: replace the specified Lease
- **/apis/coordination.k8s.io/v1/watch/namespaces/{namespace}/leases/{name}**
  - **GET**: watch changes to an object of kind Lease. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 13.2.2.1. /apis/coordination.k8s.io/v1/leases

Table 13.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Lease

**Table 13.2. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">LeaseList</a> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 13.2.2.2. /apis/coordination.k8s.io/v1/watch/leases

Table 13.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Lease. deprecated: use the 'watch' parameter with a list operation instead.

**Table 13.4. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 13.2.2.3. /apis/coordination.k8s.io/v1/namespaces/{namespace}/leases

Table 13.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 13.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Lease

Table 13.7. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 13.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 13.9. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Lease

Table 13.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 13.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">LeaseList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a Lease

**Table 13.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 13.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Lease</b> schema	

**Table 13.14. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Lease</b> schema

HTTP code	Response body
201 - Created	<a href="#">Lease</a> schema
202 - Accepted	<a href="#">Lease</a> schema
401 - Unauthorized	Empty

#### 13.2.2.4. /apis/coordination.k8s.io/v1/watch/namespaces/{namespace}/leases

Table 13.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 13.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Lease. deprecated: use the 'watch' parameter with a list operation instead.

Table 13.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 13.2.2.5. /apis/coordination.k8s.io/v1/namespaces/{namespace}/leases/{name}

Table 13.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Lease
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 13.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a Lease

Table 13.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 13.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 13.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Lease

Table 13.23. HTTP responses

HTTP code	Response body
200 - OK	<b>Lease</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Lease

**Table 13.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 13.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 13.26. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Lease</b> schema
201 - Created	<b>Lease</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Lease

Table 13.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 13.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Lease</b> schema	

Table 13.29. HTTP responses

HTTP code	Response body
200 - OK	<b>Lease</b> schema
201 - Created	<b>Lease</b> schema
401 - Unauthorized	Empty

### 13.2.2.6. /apis/coordination.k8s.io/v1/watch/namespaces/{namespace}/leases/{name}

Table 13.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Lease
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 13.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Lease. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 13.32. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty



## CHAPTER 14. CORE APIS

### 14.1. CORE APIS

#### 14.1.1. Binding [v1]

##### Description

Binding ties one object to another; for example, a pod is bound to a node by a scheduler. Deprecated in 1.7, please use the bindings subresource of pods instead.

##### Type

**object**

#### 14.1.2. ComponentStatus [v1]

##### Description

ComponentStatus (and ComponentStatusList) holds the cluster validation info. Deprecated: This API is deprecated in v1.19+

##### Type

**object**

#### 14.1.3. ConfigMap [v1]

##### Description

ConfigMap holds configuration data for pods to consume.

##### Type

**object**

#### 14.1.4. Endpoints [v1]

##### Description

Endpoints is a collection of endpoints that implement the actual service. Example:

```
Name: "mysvc",
Subsets: [
  {
    Addresses: [{"ip": "10.10.1.1"}, {"ip": "10.10.2.2"}],
    Ports: [{"name": "a", "port": 8675}, {"name": "b", "port": 309}]
  },
  {
    Addresses: [{"ip": "10.10.3.3"}],
    Ports: [{"name": "a", "port": 93}, {"name": "b", "port": 76}]
  },
]
```

##### Type

**object**

#### 14.1.5. Event [v1]

**Description**

Event is a report of an event somewhere in the cluster. Events have a limited retention time and triggers and messages may evolve with time. Event consumers should not rely on the timing of an event with a given Reason reflecting a consistent underlying trigger, or the continued existence of events with that Reason. Events should be treated as informative, best-effort, supplemental data.

**Type**

**object**

**14.1.6. LimitRange [v1]****Description**

LimitRange sets resource usage limits for each kind of resource in a Namespace.

**Type**

**object**

**14.1.7. Namespace [v1]****Description**

Namespace provides a scope for Names. Use of multiple namespaces is optional.

**Type**

**object**

**14.1.8. Node [v1]****Description**

Node is a worker node in Kubernetes. Each node will have a unique identifier in the cache (i.e. in etcd).

**Type**

**object**

**14.1.9. PersistentVolume [v1]****Description**

PersistentVolume (PV) is a storage resource provisioned by an administrator. It is analogous to a node. More info: <https://kubernetes.io/docs/concepts/storage/persistent-volumes>

**Type**

**object**

**14.1.10. PersistentVolumeClaim [v1]****Description**

PersistentVolumeClaim is a user's request for and claim to a persistent volume

**Type**

**object**

**14.1.11. Pod [v1]****Description**

Pod is a collection of containers that can run on a host. This resource is created by clients and scheduled onto hosts.

**Type****object****14.1.12. PodTemplate [v1]****Description**

PodTemplate describes a template for creating copies of a predefined pod.

**Type****object****14.1.13. ReplicationController [v1]****Description**

ReplicationController represents the configuration of a replication controller.

**Type****object****14.1.14. ResourceQuota [v1]****Description**

ResourceQuota sets aggregate quota restrictions enforced per namespace

**Type****object****14.1.15. Secret [v1]****Description**

Secret holds secret data of a certain type. The total bytes of the values in the Data field must be less than MaxSecretSize bytes.

**Type****object****14.1.16. Service [v1]****Description**

Service is a named abstraction of software service (for example, mysql) consisting of local port (for example 3306) that the proxy listens on, and the selector that determines which pods will answer requests sent through the proxy.

**Type****object****14.1.17. ServiceAccount [v1]****Description**

ServiceAccount binds together: \* a name, understood by users, and perhaps by peripheral systems, for an identity \* a principal that can be authenticated and authorized \* a set of secrets

**Type****object**

## 14.2. BINDING [V1]

**Description**

Binding ties one object to another; for example, a pod is bound to a node by a scheduler. Deprecated in 1.7, please use the bindings subresource of pods instead.

**Type****object****Required**

- **target**

### 14.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>target</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.

#### 14.2.1.1. .target

##### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

##### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

## 14.2.2. API endpoints

The following API endpoints are available:

- **/api/v1/namespaces/{namespace}/bindings**
  - **POST**: create a Binding
- **/api/v1/namespaces/{namespace}/pods/{name}/binding**
  - **POST**: create binding of a Pod

### 14.2.2.1. /api/v1/namespaces/{namespace}/bindings

Table 14.1. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.2. Global query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**POST**

Description

create a Binding

Table 14.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Binding</b> schema	

Table 14.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Binding</b> schema
201 - Created	<b>Binding</b> schema
202 - Accepted	<b>Binding</b> schema
401 - Unauthorized	Empty

#### 14.2.2.2. /api/v1/namespaces/{namespace}/pods/{name}/binding

Table 14.5. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Binding
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.6. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create binding of a Pod

Table 14.7. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Binding</b> schema	

Table 14.8. HTTP responses

HTTP code	Response body
200 - OK	<b>Binding</b> schema
201 - Created	<b>Binding</b> schema
202 - Accepted	<b>Binding</b> schema
401 - Unauthorized	Empty

**14.3. COMPONENTSTATUS [V1]**

## Description

ComponentStatus (and ComponentStatusList) holds the cluster validation info. Deprecated: This API is deprecated in v1.19+

## Type

**object**

### 14.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>conditions</b>	<b>array</b>	List of component conditions observed
<b>conditions[]</b>	<b>object</b>	Information about the condition of a component.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 14.3.1.1. .conditions

## Description

List of component conditions observed

Type

**array**

### 14.3.1.2. .conditions[]

Description

Information about the condition of a component.

Type

**object**

Required

- **type**
- **status**

Property	Type	Description
<b>error</b>	<b>string</b>	Condition error code for a component. For example, a health check error code.
<b>message</b>	<b>string</b>	Message about the condition for a component. For example, information about a health check.
<b>status</b>	<b>string</b>	Status of the condition for a component. Valid values for "Healthy": "True", "False", or "Unknown".
<b>type</b>	<b>string</b>	Type of condition for a component. Valid value: "Healthy"

### 14.3.2. API endpoints

The following API endpoints are available:

- **/api/v1/componentstatuses**
  - **GET**: list objects of kind ComponentStatus
- **/api/v1/componentstatuses/{name}**
  - **GET**: read the specified ComponentStatus

#### 14.3.2.1. /api/v1/componentstatuses

Table 14.9. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list objects of kind ComponentStatus

Table 14.10. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">ComponentStatusList</a> schema
401 - Unauthorized	Empty

### 14.3.2.2. /api/v1/componentstatuses/{name}

Table 14.11. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ComponentStatus

Table 14.12. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **GET**

#### Description

read the specified ComponentStatus

Table 14.13. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">ComponentStatus</a> schema
401 - Unauthorized	Empty

## 14.4. CONFIGMAP [V1]

#### Description

ConfigMap holds configuration data for pods to consume.

#### Type

**object**

### 14.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>binaryData</b>	<b>object (string)</b>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '.'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
<b>data</b>	<b>object (string)</b>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
<b>immutable</b>	<b>boolean</b>	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil.



Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

#### 14.4.2. API endpoints

The following API endpoints are available:

- **/api/v1/configmaps**
  - **GET**: list or watch objects of kind ConfigMap
- **/api/v1/watch/configmaps**
  - **GET**: watch individual changes to a list of ConfigMap. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/configmaps**
  - **DELETE**: delete collection of ConfigMap
  - **GET**: list or watch objects of kind ConfigMap
  - **POST**: create a ConfigMap
- **/api/v1/watch/namespaces/{namespace}/configmaps**
  - **GET**: watch individual changes to a list of ConfigMap. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/configmaps/{name}**
  - **DELETE**: delete a ConfigMap
  - **GET**: read the specified ConfigMap

- **PATCH**: partially update the specified ConfigMap
- **PUT**: replace the specified ConfigMap
- **/api/v1/watch/namespaces/{namespace}/configmaps/{name}**
  - **GET**: watch changes to an object of kind ConfigMap. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 14.4.2.1. /api/v1/configmaps

Table 14.14. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind ConfigMap

**Table 14.15. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">ConfigMapList</a> schema

HTTP code	Response body
401 - Unauthorized	Empty

#### 14.4.2.2. /api/v1/watch/configmaps

Table 14.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ConfigMap. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.4.2.3. /api/v1/namespaces/{namespace}/configmaps

Table 14.18. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ConfigMap

Table 14.20. Query parameters

Parameter	Type	Description
-----------	------	-------------



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ConfigMap

Table 14.23. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.24. HTTP responses

HTTP code	Response body
200 - OK	<b>ConfigMapList</b> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a ConfigMap

**Table 14.25. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.26. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ConfigMap</b> schema	

**Table 14.27. HTTP responses**

HTTP code	Response body
200 - OK	<b>ConfigMap</b> schema



HTTP code	Response body
201 - Created	<a href="#">ConfigMap</a> schema
202 - Accepted	<a href="#">ConfigMap</a> schema
401 - Unauthorized	Empty

#### 14.4.2.4. /api/v1/watch/namespaces/{namespace}/configmaps

Table 14.28. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.29. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ConfigMap. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.30. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.4.2.5. /api/v1/namespaces/{namespace}/configmaps/{name}

Table 14.31. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ConfigMap
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.32. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a ConfigMap

Table 14.33. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.34. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.35. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ConfigMap

Table 14.36. HTTP responses

HTTP code	Response body
200 - OK	<b>ConfigMap</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ConfigMap

**Table 14.37. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.38. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.39. HTTP responses

HTTP code	Response body
200 - OK	<b>ConfigMap</b> schema
201 - Created	<b>ConfigMap</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ConfigMap

Table 14.40. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.41. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ConfigMap</b> schema	

Table 14.42. HTTP responses

HTTP code	Response body
200 - OK	<b>ConfigMap</b> schema
201 - Created	<b>ConfigMap</b> schema
401 - Unauthorized	Empty

#### 14.4.2.6. /api/v1/watch/namespaces/{namespace}/configmaps/{name}

Table 14.43. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ConfigMap
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.44. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ConfigMap. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.45. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 14.5. ENDPOINTS [V1]

### Description

Endpoints is a collection of endpoints that implement the actual service. Example:

```
Name: "mysvc",
Subsets: [
  {
    Addresses: [{"ip": "10.10.1.1"}, {"ip": "10.10.2.2"}],
    Ports: [{"name": "a", "port": 8675}, {"name": "b", "port": 309}]
  },
  {
    Addresses: [{"ip": "10.10.3.3"}],
    Ports: [{"name": "a", "port": 93}, {"name": "b", "port": 76}]
  },
]
```

### Type

**object**

### 14.5.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>subsets</b>	<b>array</b>	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

Property	Type	Description
<b>subsets[]</b>	<b>object</b>	<p>EndpointSubset is a group of addresses with a common set of ports. The expanded set of endpoints is the Cartesian product of Addresses x Ports. For example, given:</p> <pre>{ Addresses: [{"ip": "10.10.1.1"}, {"ip": "10.10.2.2"}], Ports: [{"name": "a", "port": 8675}, {"name": "b", "port": 309}] }</pre> <p>The resulting set of endpoints can be viewed as:</p> <pre>a: [ 10.10.1.1:8675, 10.10.2.2:8675 ], b: [ 10.10.1.1:309, 10.10.2.2:309 ]</pre>

### 14.5.1.1. .subsets

#### Description

The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

#### Type

**array**

### 14.5.1.2. .subsets[]

#### Description

EndpointSubset is a group of addresses with a common set of ports. The expanded set of endpoints is the Cartesian product of Addresses x Ports. For example, given:

```
{
  Addresses: [{"ip": "10.10.1.1"}, {"ip": "10.10.2.2"}],
  Ports:    [{"name": "a", "port": 8675}, {"name": "b", "port": 309}]
}
```

The resulting set of endpoints can be viewed as:

```
a: [ 10.10.1.1:8675, 10.10.2.2:8675 ],
b: [ 10.10.1.1:309, 10.10.2.2:309 ]
```

#### Type

**object**

Property	Type	Description
<b>addresses</b>	<b>array</b>	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.
<b>addresses[]</b>	<b>object</b>	EndpointAddress is a tuple that describes single IP address.
<b>notReadyAddresses</b>	<b>array</b>	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
<b>notReadyAddresses[]</b>	<b>object</b>	EndpointAddress is a tuple that describes single IP address.
<b>ports</b>	<b>array</b>	Port numbers available on the related IP addresses.
<b>ports[]</b>	<b>object</b>	EndpointPort is a tuple that describes a single port.

#### 14.5.1.3. .subsets[].addresses

##### Description

IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

##### Type

**array**

#### 14.5.1.4. .subsets[].addresses[]

##### Description

EndpointAddress is a tuple that describes single IP address.

##### Type

**object**

##### Required

- **ip**



Property	Type	Description
<b>hostname</b>	<b>string</b>	The Hostname of this endpoint
<b>ip</b>	<b>string</b>	The IP of this endpoint. May not be loopback (127.0.0.0/8 or ::1), link-local (169.254.0.0/16 or fe80::/10), or link-local multicast (224.0.0.0/24 or ff02::/16).
<b>nodeName</b>	<b>string</b>	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
<b>targetRef</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.

#### 14.5.1.5. .subsets[].addresses[].targetRef

##### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

##### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

### 14.5.1.6. .subsets[].notReadyAddresses

#### Description

IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.

#### Type

**array**

### 14.5.1.7. .subsets[].notReadyAddresses[]

#### Description

EndpointAddress is a tuple that describes single IP address.

#### Type

**object**

#### Required

- **ip**

Property	Type	Description
<b>hostname</b>	<b>string</b>	The Hostname of this endpoint
<b>ip</b>	<b>string</b>	The IP of this endpoint. May not be loopback (127.0.0.0/8 or ::1), link-local (169.254.0.0/16 or fe80::/10), or link-local multicast (224.0.0.0/24 or ff02::/16).
<b>nodeName</b>	<b>string</b>	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
<b>targetRef</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.

### 14.5.1.8. .subsets[].notReadyAddresses[].targetRef

#### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

#### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>

Property	Type	Description
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

#### 14.5.1.9. .subsets[].ports

##### Description

Port numbers available on the related IP addresses.

##### Type

**array**

#### 14.5.1.10. .subsets[].ports[]

##### Description

EndpointPort is a tuple that describes a single port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>appProtocol</b>	<b>string</b>	<p>The application protocol for this port. This is used as a hint for implementations to offer richer behavior for protocols that they understand. This field follows standard Kubernetes label syntax. Valid values are either:</p> <ul style="list-style-type: none"> <li>* Un-prefixed protocol names - reserved for IANA standard service names (as per RFC-6335 and <a href="https://www.iana.org/assignments/service-names">https://www.iana.org/assignments/service-names</a>).</li> <li>* Kubernetes-defined prefixed names: * 'kubernetes.io/h2c' - HTTP/2 over cleartext as described in <a href="https://www.rfc-editor.org/rfc/rfc7540">https://www.rfc-editor.org/rfc/rfc7540</a></li> <li>* Other protocols should use implementation-defined prefixed names such as mycompany.com/my-custom-protocol.</li> </ul>
<b>name</b>	<b>string</b>	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
<b>port</b>	<b>integer</b>	The port number of the endpoint.
<b>protocol</b>	<b>string</b>	<p>The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.</p> <p>Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.</p>

### 14.5.2. API endpoints

The following API endpoints are available:

- **/api/v1/endpoints**
  - **GET**: list or watch objects of kind Endpoints

- **/api/v1/watch/endpoints**
  - **GET**: watch individual changes to a list of Endpoints. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/endpoints**
  - **DELETE**: delete collection of Endpoints
  - **GET**: list or watch objects of kind Endpoints
  - **POST**: create Endpoints
- **/api/v1/watch/namespaces/{namespace}/endpoints**
  - **GET**: watch individual changes to a list of Endpoints. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/endpoints/{name}**
  - **DELETE**: delete Endpoints
  - **GET**: read the specified Endpoints
  - **PATCH**: partially update the specified Endpoints
  - **PUT**: replace the specified Endpoints
- **/api/v1/watch/namespaces/{namespace}/endpoints/{name}**
  - **GET**: watch changes to an object of kind Endpoints. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 14.5.2.1. /api/v1/endpoints

Table 14.46. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the <code>continue</code> field to determine whether more results are available. Servers may choose not to support the <code>limit</code> argument and will return all of the available results. If <code>limit</code> is specified and the <code>continue</code> field is empty, clients may assume that no more results are available. This field is not supported if <code>watch</code> is true.</p> <p>The server guarantees that the objects returned when using <code>continue</code> will be identical to issuing a single list call without a <code>limit</code> - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using <code>limit</code> to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p><code>resourceVersion</code> sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p><code>resourceVersionMatch</code> determines how <code>resourceVersion</code> is applied to list calls. It is highly recommended that <code>resourceVersionMatch</code> be set for list calls where <code>resourceVersion</code> is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Endpoints

**Table 14.47. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">EndpointsList</a> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 14.5.2.2. /api/v1/watch/endpoints

Table 14.48. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Endpoints. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.49. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.5.2.3. /api/v1/namespaces/{namespace}/endpoints

Table 14.50. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.51. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Endpoints

Table 14.52. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.53. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.54. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Endpoints

Table 14.55. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.56. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">EndpointsList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create Endpoints

**Table 14.57. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.58. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Endpoints</b> schema	

**Table 14.59. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Endpoints</b> schema

HTTP code	Response body
201 - Created	<a href="#">Endpoints</a> schema
202 - Accepted	<a href="#">Endpoints</a> schema
401 - Unauthorized	Empty

#### 14.5.2.4. /api/v1/watch/namespaces/{namespace}/endpoints

Table 14.60. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.61. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Endpoints. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.62. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.5.2.5. /api/v1/namespaces/{namespace}/endpoints/{name}

Table 14.63. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Endpoints
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.64. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete Endpoints

Table 14.65. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.66. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.67. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Endpoints

Table 14.68. HTTP responses

HTTP code	Response body
200 - OK	<b>Endpoints</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Endpoints

**Table 14.69. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.70. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.71. HTTP responses

HTTP code	Response body
200 - OK	<b>Endpoints</b> schema
201 - Created	<b>Endpoints</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Endpoints

Table 14.72. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.73. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Endpoints</b> schema	

Table 14.74. HTTP responses

HTTP code	Response body
200 - OK	<b>Endpoints</b> schema
201 - Created	<b>Endpoints</b> schema
401 - Unauthorized	Empty

#### 14.5.2.6. /api/v1/watch/namespaces/{namespace}/endpoints/{name}

Table 14.75. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Endpoints
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.76. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Endpoints. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.77. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>WatchEvent</b> schema
401 - Unauthorized	Empty

## 14.6. EVENT [V1]

### Description

Event is a report of an event somewhere in the cluster. Events have a limited retention time and triggers and messages may evolve with time. Event consumers should not rely on the timing of an event with a given Reason reflecting a consistent underlying trigger, or the continued existence of events with that Reason. Events should be treated as informative, best-effort, supplemental data.

### Type

**object**

### Required

- **metadata**
- **involvedObject**

### 14.6.1. Specification

Property	Type	Description
<b>action</b>	<b>string</b>	What action was taken/failed regarding to the Regarding object.
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>count</b>	<b>integer</b>	The number of times this event has occurred.
<b>eventTime</b>	<b>MicroTime</b>	Time when this Event was first observed.

Property	Type	Description
<b>firstTimestamp</b>	<b>Time</b>	The time at which the event was first recorded. (Time of server receipt is in TypeMeta.)
<b>involvedObject</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>lastTimestamp</b>	<b>Time</b>	The time at which the most recent occurrence of this event was recorded.
<b>message</b>	<b>string</b>	A human-readable description of the status of this operation.
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>reason</b>	<b>string</b>	This should be a short, machine understandable string that gives the reason for the transition into the object's current status.
<b>related</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.
<b>reportingComponent</b>	<b>string</b>	Name of the controller that emitted this Event, e.g. <b>kubernetes.io/kubelet</b> .

Property	Type	Description
<b>reportingInstance</b>	<b>string</b>	ID of the controller instance, e.g. <b>kubelet-xyzf</b> .
<b>series</b>	<b>object</b>	EventSeries contain information on series of events, i.e. thing that was/is happening continuously for some time.
<b>source</b>	<b>object</b>	EventSource contains information for an event.
<b>type</b>	<b>string</b>	Type of this event (Normal, Warning), new types could be added in the future

#### 14.6.1.1. `.involvedObject`

##### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

##### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

#### 14.6.1.2. .related

##### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

##### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

### 14.6.1.3. .series

#### Description

EventSeries contain information on series of events, i.e. thing that was/is happening continuously for some time.

#### Type

**object**

Property	Type	Description
<b>count</b>	<b>integer</b>	Number of occurrences in this series up to the last heartbeat time
<b>lastObservedTime</b>	<b>MicroTime</b>	Time of the last occurrence observed

### 14.6.1.4. .source

#### Description

EventSource contains information for an event.

#### Type

**object**

Property	Type	Description
<b>component</b>	<b>string</b>	Component from which the event is generated.
<b>host</b>	<b>string</b>	Node name on which the event is generated.

## 14.6.2. API endpoints

The following API endpoints are available:

- **/api/v1/events**
  - **GET**: list or watch objects of kind Event
- **/api/v1/watch/events**
  - **GET**: watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/events**
  - **DELETE**: delete collection of Event
  - **GET**: list or watch objects of kind Event



- **POST**: create an Event
- **/api/v1/watch/namespaces/{namespace}/events**
  - **GET**: watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/events/{name}**
  - **DELETE**: delete an Event
  - **GET**: read the specified Event
  - **PATCH**: partially update the specified Event
  - **PUT**: replace the specified Event
- **/api/v1/watch/namespaces/{namespace}/events/{name}**
  - **GET**: watch changes to an object of kind Event. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 14.6.2.1. /api/v1/events

Table 14.78. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Event

**Table 14.79. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>EventList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 14.6.2.2. /api/v1/watch/events

Table 14.80. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.81. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.6.2.3. /api/v1/namespaces/{namespace}/events

Table 14.82. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.83. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Event

Table 14.84. Query parameters

Parameter	Type	Description
-----------	------	-------------



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.85. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.86. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Event

Table 14.87. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.88. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">EventList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create an Event

**Table 14.89. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.90. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Event</b> schema	

**Table 14.91. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Event</b> schema



HTTP code	Response body
201 - Created	<b>Event</b> schema
202 - Accepted	<b>Event</b> schema
401 - Unauthorized	Empty

#### 14.6.2.4. /api/v1/watch/namespaces/{namespace}/events

Table 14.92. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.93. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.94. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.6.2.5. /api/v1/namespaces/{namespace}/events/{name}

Table 14.95. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Event
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.96. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete an Event

Table 14.97. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.98. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.99. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Event

Table 14.100. HTTP responses

HTTP code	Response body
200 - OK	<b>Event</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Event

**Table 14.101. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.102. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.103. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Event</b> schema
201 - Created	<b>Event</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Event

Table 14.104. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.105. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Event</b> schema	

Table 14.106. HTTP responses

HTTP code	Response body
200 - OK	<b>Event</b> schema
201 - Created	<b>Event</b> schema
401 - Unauthorized	Empty

#### 14.6.2.6. /api/v1/watch/namespaces/{namespace}/events/{name}

Table 14.107. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Event
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.108. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Event. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.109. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 14.7. LIMITRANGE [V1]

### Description

LimitRange sets resource usage limits for each kind of resource in a Namespace.

### Type

**object**

### 14.7.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>spec</b>	<b>object</b>	LimitRangeSpec defines a min/max usage limit for resources that match on kind.

#### 14.7.1.1. .spec

##### Description

LimitRangeSpec defines a min/max usage limit for resources that match on kind.

##### Type

**object**

##### Required

- **limits**

Property	Type	Description
<b>limits</b>	<b>array</b>	Limits is the list of LimitRangeItem objects that are enforced.
<b>limits[]</b>	<b>object</b>	LimitRangeItem defines a min/max usage limit for any resource that matches on kind.

#### 14.7.1.2. .spec.limits

##### Description

Limits is the list of LimitRangeItem objects that are enforced.

##### Type

**array**

#### 14.7.1.3. .spec.limits[]

##### Description

LimitRangeItem defines a min/max usage limit for any resource that matches on kind.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
<b>default</b>	<b>object (Quantity)</b>	Default resource requirement limit value by resource name if resource limit is omitted.
<b>defaultRequest</b>	<b>object (Quantity)</b>	DefaultRequest is the default resource requirement request value by resource name if resource request is omitted.
<b>max</b>	<b>object (Quantity)</b>	Max usage constraints on this kind by resource name.
<b>maxLimitRequestRatio</b>	<b>object (Quantity)</b>	MaxLimitRequestRatio if specified, the named resource must have a request and limit that are both non-zero where limit divided by request is less than or equal to the enumerated value; this represents the max burst for the named resource.
<b>min</b>	<b>object (Quantity)</b>	Min usage constraints on this kind by resource name.
<b>type</b>	<b>string</b>	Type of resource that this limit applies to.

### 14.7.2. API endpoints

The following API endpoints are available:

- **/api/v1/limitranges**
  - **GET**: list or watch objects of kind LimitRange
- **/api/v1/watch/limitranges**
  - **GET**: watch individual changes to a list of LimitRange. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/limitranges**
  - **DELETE**: delete collection of LimitRange
  - **GET**: list or watch objects of kind LimitRange
  - **POST**: create a LimitRange
- **/api/v1/watch/namespaces/{namespace}/limitranges**
  - **GET**: watch individual changes to a list of LimitRange. deprecated: use the 'watch' parameter with a list operation instead.

- **/api/v1/namespaces/{namespace}/limitranges/{name}**
  - **DELETE:** delete a LimitRange
  - **GET:** read the specified LimitRange
  - **PATCH:** partially update the specified LimitRange
  - **PUT:** replace the specified LimitRange
- **/api/v1/watch/namespaces/{namespace}/limitranges/{name}**
  - **GET:** watch changes to an object of kind LimitRange. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 14.7.2.1. /api/v1/limitranges

Table 14.110. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind LimitRange

Table 14.111. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">LimitRangeList</a> schema
401 - Unauthorized	Empty

### 14.7.2.2. /api/v1/watch/limitranges

Table 14.112. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of LimitRange. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.113. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.7.2.3. /api/v1/namespaces/{namespace}/limitranges

Table 14.114. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.115. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of LimitRange

Table 14.116. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.117. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.118. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind LimitRange

Table 14.119. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the <code>continue</code> field to determine whether more results are available. Servers may choose not to support the <code>limit</code> argument and will return all of the available results. If <code>limit</code> is specified and the <code>continue</code> field is empty, clients may assume that no more results are available. This field is not supported if <code>watch</code> is true.</p> <p>The server guarantees that the objects returned when using <code>continue</code> will be identical to issuing a single list call without a <code>limit</code> - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using <code>limit</code> to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p><code>resourceVersion</code> sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p><code>resourceVersionMatch</code> determines how <code>resourceVersion</code> is applied to list calls. It is highly recommended that <code>resourceVersionMatch</code> be set for list calls where <code>resourceVersion</code> is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.120. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">LimitRangeList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a `LimitRange`

**Table 14.121. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized <code>dryRun</code> directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	<code>fieldManager</code> is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	<code>fieldValidation</code> instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a <code>BadRequest</code> error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.122. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>LimitRange</b> schema	

**Table 14.123. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>LimitRange</b> schema

HTTP code	Response body
201 - Created	<a href="#">LimitRange</a> schema
202 - Accepted	<a href="#">LimitRange</a> schema
401 - Unauthorized	Empty

#### 14.7.2.4. /api/v1/watch/namespaces/{namespace}/limitranges

Table 14.124. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.125. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of LimitRange. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.126. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.7.2.5. /api/v1/namespaces/{namespace}/limitranges/{name}

Table 14.127. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the LimitRange
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.128. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a LimitRange

Table 14.129. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.130. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.131. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified LimitRange

Table 14.132. HTTP responses

HTTP code	Response body
200 - OK	<b>LimitRange</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified LimitRange

**Table 14.133. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.134. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.135. HTTP responses

HTTP code	Reponse body
200 - OK	<b>LimitRange</b> schema
201 - Created	<b>LimitRange</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified LimitRange

Table 14.136. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.137. Body parameters

Parameter	Type	Description
<b>body</b>	<b>LimitRange</b> schema	

Table 14.138. HTTP responses

HTTP code	Response body
200 - OK	<b>LimitRange</b> schema
201 - Created	<b>LimitRange</b> schema
401 - Unauthorized	Empty

#### 14.7.2.6. /api/v1/watch/namespaces/{namespace}/limitranges/{name}

Table 14.139. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the LimitRange
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects



Table 14.140. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind LimitRange. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.141. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 14.8. NAMESPACE [V1]

### Description

Namespace provides a scope for Names. Use of multiple namespaces is optional.

### Type

**object**

### 14.8.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	NamespaceSpec describes the attributes on a Namespace.

Property	Type	Description
<b>status</b>	<b>object</b>	NamespaceStatus is information about the current status of a Namespace.

#### 14.8.1.1. .spec

##### Description

NamespaceSpec describes the attributes on a Namespace.

##### Type

**object**

Property	Type	Description
<b>finalizers</b>	<b>array (string)</b>	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: <a href="https://kubernetes.io/docs/tasks/administer-cluster/namespaces/">https://kubernetes.io/docs/tasks/administer-cluster/namespaces/</a>

#### 14.8.1.2. .status

##### Description

NamespaceStatus is information about the current status of a Namespace.

##### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array</b>	Represents the latest available observations of a namespace's current state.
<b>conditions[]</b>	<b>object</b>	NamespaceCondition contains details about state of namespace.

Property	Type	Description
<b>phase</b>	<b>string</b>	<p>Phase is the current lifecycle phase of the namespace. More info: <a href="https://kubernetes.io/docs/tasks/administer-cluster/namespaces/">https://kubernetes.io/docs/tasks/administer-cluster/namespaces/</a></p> <p>Possible enum values: - <b>"Active"</b> means the namespace is available for use in the system - <b>"Terminating"</b> means the namespace is undergoing graceful termination</p>

#### 14.8.1.3. .status.conditions

##### Description

Represents the latest available observations of a namespace's current state.

##### Type

**array**

#### 14.8.1.4. .status.conditions[]

##### Description

NamespaceCondition contains details about state of namespace.

##### Type

**object**

##### Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	
<b>message</b>	<b>string</b>	
<b>reason</b>	<b>string</b>	
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.

Property	Type	Description
<b>type</b>	<b>string</b>	Type of namespace controller condition.

## 14.8.2. API endpoints

The following API endpoints are available:

- **/api/v1/namespaces**
  - **GET**: list or watch objects of kind Namespace
  - **POST**: create a Namespace
- **/api/v1/watch/namespaces**
  - **GET**: watch individual changes to a list of Namespace. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{name}**
  - **DELETE**: delete a Namespace
  - **GET**: read the specified Namespace
  - **PATCH**: partially update the specified Namespace
  - **PUT**: replace the specified Namespace
- **/api/v1/watch/namespaces/{name}**
  - **GET**: watch changes to an object of kind Namespace. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/namespaces/{name}/status**
  - **GET**: read status of the specified Namespace
  - **PATCH**: partially update status of the specified Namespace
  - **PUT**: replace status of the specified Namespace
- **/api/v1/namespaces/{name}/finalize**
  - **PUT**: replace finalize of the specified Namespace

### 14.8.2.1. /api/v1/namespaces

Table 14.142. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

list or watch objects of kind Namespace

Table 14.143. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.144. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">NamespaceList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a Namespace

**Table 14.145. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.146. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Namespace</b> schema	

**Table 14.147. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Namespace</b> schema

HTTP code	Response body
201 - Created	<a href="#">Namespace</a> schema
202 - Accepted	<a href="#">Namespace</a> schema
401 - Unauthorized	Empty

### 14.8.2.2. /api/v1/watch/namespaces

Table 14.148. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
**GET**

**Description**

watch individual changes to a list of Namespace. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.149. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**14.8.2.3. /api/v1/namespaces/{name}****Table 14.150. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Namespace

**Table 14.151. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a Namespace

**Table 14.152. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.153. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.154. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Namespace

Table 14.155. HTTP responses

HTTP code	Response body
200 - OK	<b>Namespace</b> schema



HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Namespace

**Table 14.156. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.157. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.158. HTTP responses

HTTP code	Response body
200 - OK	<b>Namespace</b> schema
201 - Created	<b>Namespace</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Namespace

Table 14.159. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.160. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Namespace</b> schema	

Table 14.161. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Namespace</b> schema
201 - Created	<b>Namespace</b> schema
401 - Unauthorized	Empty

#### 14.8.2.4. /api/v1/watch/namespaces/{name}

Table 14.162. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Namespace

Table 14.163. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Namespace. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.164. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.8.2.5. /api/v1/namespaces/{name}/status

Table 14.165. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Namespace

Table 14.166. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Namespace

Table 14.167. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">Namespace</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Namespace

Table 14.168. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.169. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.170. HTTP responses



HTTP code	Response body
200 - OK	<b>Namespace</b> schema
201 - Created	<b>Namespace</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Namespace

**Table 14.171. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.172. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Namespace</b> schema	

Table 14.173. HTTP responses

HTTP code	Response body
200 - OK	<b>Namespace</b> schema
201 - Created	<b>Namespace</b> schema
401 - Unauthorized	Empty

#### 14.8.2.6. /api/v1/namespaces/{name}/finalize

Table 14.174. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Namespace

Table 14.175. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****PUT****Description**

replace finalize of the specified Namespace

**Table 14.176. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Namespace</b> schema	

**Table 14.177. HTTP responses**

HTTP code	Response body
200 - OK	<b>Namespace</b> schema
201 - Created	<b>Namespace</b> schema
401 - Unauthorized	Empty

**14.9. NODE [V1]****Description**

Node is a worker node in Kubernetes. Each node will have a unique identifier in the cache (i.e. in etcd).

## Type

**object**

### 14.9.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	NodeSpec describes the attributes that a node is created with.
<b>status</b>	<b>object</b>	NodeStatus is information about the current status of a node.

#### 14.9.1.1. .spec

##### Description

NodeSpec describes the attributes that a node is created with.

##### Type

**object**

Property	Type	Description
<b>configSource</b>	<b>object</b>	NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22
<b>externalID</b>	<b>string</b>	Deprecated. Not all kubelets will set this field. Remove field after 1.13. see: <a href="https://issues.k8s.io/61966">https://issues.k8s.io/61966</a>
<b>podCIDR</b>	<b>string</b>	PodCIDR represents the pod IP range assigned to the node.
<b>podCIDRs</b>	<b>array (string)</b>	podCIDRs represents the IP ranges assigned to the node for usage by Pods on that node. If this field is specified, the 0th entry must match the podCIDR field. It may contain at most 1 value for each of IPv4 and IPv6.
<b>providerID</b>	<b>string</b>	ID of the node assigned by the cloud provider in the format: <ProviderName>://<ProviderSpecificNodeID>
<b>taints</b>	<b>array</b>	If specified, the node's taints.
<b>taints[]</b>	<b>object</b>	The node this Taint is attached to has the "effect" on any pod that does not tolerate the Taint.
<b>unschedulable</b>	<b>boolean</b>	Unschedulable controls node schedulability of new pods. By default, node is schedulable. More info: <a href="https://kubernetes.io/docs/concepts/nodes/node/#manual-node-administration">https://kubernetes.io/docs/concepts/nodes/node/#manual-node-administration</a>

**14.9.1.2. .spec.configSource****Description**

NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22

**Type**

**object**

Property	Type	Description
<b>configMap</b>	<b>object</b>	ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <a href="https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration">https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration</a>

**14.9.1.3. .spec.configSource.configMap****Description**

ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration>

**Type****object****Required**

- **namespace**
- **name**
- **kubeletConfigKey**

Property	Type	Description
<b>kubeletConfigKey</b>	<b>string</b>	KubeletConfigKey declares which key of the referenced ConfigMap corresponds to the KubeletConfiguration structure. This field is required in all cases.
<b>name</b>	<b>string</b>	Name is the metadata.name of the referenced ConfigMap. This field is required in all cases.
<b>namespace</b>	<b>string</b>	Namespace is the metadata.namespace of the referenced ConfigMap. This field is required in all cases.

Property	Type	Description
<b>resourceVersion</b>	<b>string</b>	ResourceVersion is the metadata.ResourceVersion of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.
<b>uid</b>	<b>string</b>	UID is the metadata.UID of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.

#### 14.9.1.4. .spec.taints

##### Description

If specified, the node's taints.

##### Type

**array**

#### 14.9.1.5. .spec.taints[]

##### Description

The node this Taint is attached to has the "effect" on any pod that does not tolerate the Taint.

##### Type

**object**

##### Required

- **key**
- **effect**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>effect</b>	<b>string</b>	<p>Required. The effect of the taint on pods that do not tolerate the taint. Valid effects are NoSchedule, PreferNoSchedule and NoExecute.</p> <p>Possible enum values: -  <b>"NoExecute"</b> Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController. -  <b>"NoSchedule"</b> Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler. -  <b>"PreferNoSchedule"</b> Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler.</p>
<b>key</b>	<b>string</b>	Required. The taint key to be applied to a node.
<b>timeAdded</b>	<b>Time</b>	TimeAdded represents the time at which the taint was added. It is only written for NoExecute taints.
<b>value</b>	<b>string</b>	The taint value corresponding to the taint key.

#### 14.9.1.6. .status

##### Description

NodeStatus is information about the current status of a node.

##### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>addresses</b>	<b>array</b>	List of addresses reachable to the node. Queried from cloud provider, if available. More info: <a href="https://kubernetes.io/docs/concepts/nodes/node/#addresses">https://kubernetes.io/docs/concepts/nodes/node/#addresses</a> Note: This field is declared as mergeable, but the merge key is not sufficiently unique, which can cause data corruption when it is merged. Callers should instead use a full-replacement patch. See <a href="https://pr.k8s.io/79391">https://pr.k8s.io/79391</a> for an example. Consumers should assume that addresses can change during the lifetime of a Node. However, there are some exceptions where this may not be possible, such as Pods that inherit a Node's address in its own status or consumers of the downward API (status.hostIP).
<b>addresses[]</b>	<b>object</b>	NodeAddress contains information for the node's address.
<b>allocatable</b>	<b>object (Quantity)</b>	Allocatable represents the resources of a node that are available for scheduling. Defaults to Capacity.
<b>capacity</b>	<b>object (Quantity)</b>	Capacity represents the total resources of a node. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#capacity">https://kubernetes.io/docs/concepts/storage/persistent-volumes#capacity</a>
<b>conditions</b>	<b>array</b>	Conditions is an array of current observed node conditions. More info: <a href="https://kubernetes.io/docs/concepts/nodes/node/#condition">https://kubernetes.io/docs/concepts/nodes/node/#condition</a>
<b>conditions[]</b>	<b>object</b>	NodeCondition contains condition information for a node.

Property	Type	Description
<b>config</b>	<b>object</b>	NodeConfigStatus describes the status of the config assigned by Node.Spec.ConfigSource.
<b>daemonEndpoints</b>	<b>object</b>	NodeDaemonEndpoints lists ports opened by daemons running on the Node.
<b>images</b>	<b>array</b>	List of container images on this node
<b>images[]</b>	<b>object</b>	Describe a container image
<b>nodeInfo</b>	<b>object</b>	NodeSystemInfo is a set of ids/uuids to uniquely identify the node.
<b>phase</b>	<b>string</b>	<p>NodePhase is the recently observed lifecycle phase of the node. More info: <a href="https://kubernetes.io/docs/concepts/nodes/node/#phase">https://kubernetes.io/docs/concepts/nodes/node/#phase</a> The field is never populated, and now is deprecated.</p> <p>Possible enum values: - <b>"Pending"</b> means the node has been created/added by the system, but not configured. - <b>"Running"</b> means the node has been configured and has Kubernetes components running. - <b>"Terminated"</b> means the node has been removed from the cluster.</p>
<b>volumesAttached</b>	<b>array</b>	List of volumes that are attached to the node.
<b>volumesAttached[]</b>	<b>object</b>	AttachedVolume describes a volume attached to a node
<b>volumesInUse</b>	<b>array (string)</b>	List of attachable volumes in use (mounted) by the node.

#### 14.9.1.7. .status.addresses

##### Description

List of addresses reachable to the node. Queried from cloud provider, if available. More info:

<https://kubernetes.io/docs/concepts/nodes/node/#addresses> Note: This field is declared as mergeable, but the merge key is not sufficiently unique, which can cause data corruption when it is merged. Callers should instead use a full-replacement patch. See <https://pr.k8s.io/79391> for an example. Consumers should assume that addresses can change during the lifetime of a Node. However, there are some exceptions where this may not be possible, such as Pods that inherit a Node's address in its own status or consumers of the downward API (status.hostIP).

**Type****array****14.9.1.8. .status.addresses[]****Description**

NodeAddress contains information for the node's address.

**Type****object****Required**

- **type**
- **address**

Property	Type	Description
<b>address</b>	<b>string</b>	The node address.
<b>type</b>	<b>string</b>	Node address type, one of Hostname, ExternallIP or InternallIP.

**14.9.1.9. .status.conditions****Description**

Conditions is an array of current observed node conditions. More info: <https://kubernetes.io/docs/concepts/nodes/node/#condition>

**Type****array****14.9.1.10. .status.conditions[]****Description**

NodeCondition contains condition information for a node.

**Type****object****Required**

- **type**
- **status**

Property	Type	Description
<b>lastHeartbeatTime</b>	<b>Time</b>	Last time we got an update on a given condition.
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transit from one status to another.
<b>message</b>	<b>string</b>	Human readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	(brief) reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of node condition.

#### 14.9.1.11. .status.config

##### Description

NodeConfigStatus describes the status of the config assigned by Node.Spec.ConfigSource.

##### Type

**object**

Property	Type	Description
<b>active</b>	<b>object</b>	NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22
<b>assigned</b>	<b>object</b>	NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22

Property	Type	Description
<b>error</b>	<b>string</b>	Error describes any problems reconciling the Spec.ConfigSource to the Active config. Errors may occur, for example, attempting to checkpoint Spec.ConfigSource to the local Assigned record, attempting to checkpoint the payload associated with Spec.ConfigSource, attempting to load or validate the Assigned config, etc. Errors may occur at different points while syncing config. Earlier errors (e.g. download or checkpointing errors) will not result in a rollback to LastKnownGood, and may resolve across Kubelet retries. Later errors (e.g. loading or validating a checkpointed config) will result in a rollback to LastKnownGood. In the latter case, it is usually possible to resolve the error by fixing the config assigned in Spec.ConfigSource. You can find additional information for debugging by searching the error message in the Kubelet log. Error is a human-readable description of the error state; machines can check whether or not Error is empty, but should not rely on the stability of the Error text across Kubelet versions.
<b>lastKnownGood</b>	<b>object</b>	NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22

#### 14.9.1.12. .status.config.active

##### Description

NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22

##### Type

**object**

Property	Type	Description
<b>configMap</b>	<b>object</b>	ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <a href="https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration">https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration</a>

#### 14.9.1.13. .status.config.active.configMap

##### Description

ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration>

##### Type

**object**

##### Required

- **namespace**
- **name**
- **kubeletConfigKey**

Property	Type	Description
<b>kubeletConfigKey</b>	<b>string</b>	KubeletConfigKey declares which key of the referenced ConfigMap corresponds to the KubeletConfiguration structure. This field is required in all cases.
<b>name</b>	<b>string</b>	Name is the metadata.name of the referenced ConfigMap. This field is required in all cases.
<b>namespace</b>	<b>string</b>	Namespace is the metadata.namespace of the referenced ConfigMap. This field is required in all cases.
<b>resourceVersion</b>	<b>string</b>	ResourceVersion is the metadata.ResourceVersion of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.

Property	Type	Description
<b>uid</b>	<b>string</b>	UID is the metadata.UID of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.

#### 14.9.1.14. .status.config.assigned

##### Description

NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22

##### Type

**object**

Property	Type	Description
<b>configMap</b>	<b>object</b>	ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <a href="https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration">https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration</a>

#### 14.9.1.15. .status.config.assigned.configMap

##### Description

ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration>

##### Type

**object**

##### Required

- **namespace**
- **name**
- **kubeletConfigKey**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>kubeletConfigKey</b>	<b>string</b>	KubeletConfigKey declares which key of the referenced ConfigMap corresponds to the KubeletConfiguration structure. This field is required in all cases.
<b>name</b>	<b>string</b>	Name is the metadata.name of the referenced ConfigMap. This field is required in all cases.
<b>namespace</b>	<b>string</b>	Namespace is the metadata.namespace of the referenced ConfigMap. This field is required in all cases.
<b>resourceVersion</b>	<b>string</b>	ResourceVersion is the metadata.ResourceVersion of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.
<b>uid</b>	<b>string</b>	UID is the metadata.UID of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.

#### 14.9.1.16. .status.config.lastKnownGood

##### Description

NodeConfigSource specifies a source of node configuration. Exactly one subfield (excluding metadata) must be non-nil. This API is deprecated since 1.22.

##### Type

**object**

Property	Type	Description
<b>configMap</b>	<b>object</b>	ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <a href="https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration">https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration</a>

#### 14.9.1.17. .status.config.lastKnownGood.configMap



**Description**

ConfigMapNodeConfigSource contains the information to reference a ConfigMap as a config source for the Node. This API is deprecated since 1.22: <https://git.k8s.io/enhancements/keps/sig-node/281-dynamic-kubelet-configuration>

**Type**

**object**

**Required**

- **namespace**
- **name**
- **kubeletConfigKey**

Property	Type	Description
<b>kubeletConfigKey</b>	<b>string</b>	KubeletConfigKey declares which key of the referenced ConfigMap corresponds to the KubeletConfiguration structure. This field is required in all cases.
<b>name</b>	<b>string</b>	Name is the metadata.name of the referenced ConfigMap. This field is required in all cases.
<b>namespace</b>	<b>string</b>	Namespace is the metadata.namespace of the referenced ConfigMap. This field is required in all cases.
<b>resourceVersion</b>	<b>string</b>	ResourceVersion is the metadata.ResourceVersion of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.
<b>uid</b>	<b>string</b>	UID is the metadata.UID of the referenced ConfigMap. This field is forbidden in Node.Spec, and required in Node.Status.

**14.9.1.18. .status.daemonEndpoints****Description**

NodeDaemonEndpoints lists ports opened by daemons running on the Node.

**Type**

**object**

Property	Type	Description
<b>kubeletEndpoint</b>	<b>object</b>	DaemonEndpoint contains information about a single Daemon endpoint.

#### 14.9.1.19. .status.daemonEndpoints.kubeletEndpoint

##### Description

DaemonEndpoint contains information about a single Daemon endpoint.

##### Type

**object**

##### Required

- **Port**

Property	Type	Description
<b>Port</b>	<b>integer</b>	Port number of the given endpoint.

#### 14.9.1.20. .status.images

##### Description

List of container images on this node

##### Type

**array**

#### 14.9.1.21. .status.images[]

##### Description

Describe a container image

##### Type

**object**

Property	Type	Description
<b>names</b>	<b>array (string)</b>	Names by which this image is known. e.g. ["kubernetes.example/hyperkube:v1.0.7", "cloud-vendor.registry.example/cloud-vendor/hyperkube:v1.0.7"]
<b>sizeBytes</b>	<b>integer</b>	The size of the image in bytes.

14.9.1.22. `.status.nodeInfo`

## Description

NodeSystemInfo is a set of ids/uuids to uniquely identify the node.

## Type

**object**

## Required

- **machineID**
- **systemUUID**
- **bootID**
- **kernelVersion**
- **osImage**
- **containerRuntimeVersion**
- **kubeletVersion**
- **kubeProxyVersion**
- **operatingSystem**
- **architecture**

Property	Type	Description
<b>architecture</b>	<b>string</b>	The Architecture reported by the node
<b>bootID</b>	<b>string</b>	Boot ID reported by the node.
<b>containerRuntimeVersion</b>	<b>string</b>	ContainerRuntime Version reported by the node through runtime remote API (e.g. containerd://1.4.2).
<b>kernelVersion</b>	<b>string</b>	Kernel Version reported by the node from 'uname -r' (e.g. 3.16.0-0.bpo.4-amd64).
<b>kubeProxyVersion</b>	<b>string</b>	KubeProxy Version reported by the node.
<b>kubeletVersion</b>	<b>string</b>	Kubelet Version reported by the node.

Property	Type	Description
<b>machineID</b>	<b>string</b>	MachineID reported by the node. For unique machine identification in the cluster this field is preferred. Learn more from man(5) machine-id: <a href="http://man7.org/linux/man-pages/man5/machine-id.5.html">http://man7.org/linux/man-pages/man5/machine-id.5.html</a>
<b>operatingSystem</b>	<b>string</b>	The Operating System reported by the node
<b>osImage</b>	<b>string</b>	OS Image reported by the node from /etc/os-release (e.g. Debian GNU/Linux 7 (wheezy)).
<b>systemUUID</b>	<b>string</b>	SystemUUID reported by the node. For unique machine identification MachineID is preferred. This field is specific to Red Hat hosts <a href="https://access.redhat.com/documentation/en-us/red_hat_subscription_management/1/html/rhsm/uuid">https://access.redhat.com/documentation/en-us/red_hat_subscription_management/1/html/rhsm/uuid</a>

#### 14.9.1.23. .status.volumesAttached

##### Description

List of volumes that are attached to the node.

##### Type

**array**

#### 14.9.1.24. .status.volumesAttached[]

##### Description

AttachedVolume describes a volume attached to a node

##### Type

**object**

##### Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	DevicePath represents the device path where the volume should be available
<b>name</b>	<b>string</b>	Name of the attached volume

## 14.9.2. API endpoints

The following API endpoints are available:

- **/api/v1/nodes**
  - **DELETE**: delete collection of Node
  - **GET**: list or watch objects of kind Node
  - **POST**: create a Node
- **/api/v1/watch/nodes**
  - **GET**: watch individual changes to a list of Node. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/nodes/{name}**
  - **DELETE**: delete a Node
  - **GET**: read the specified Node
  - **PATCH**: partially update the specified Node
  - **PUT**: replace the specified Node
- **/api/v1/watch/nodes/{name}**
  - **GET**: watch changes to an object of kind Node. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/nodes/{name}/status**
  - **GET**: read status of the specified Node
  - **PATCH**: partially update status of the specified Node
  - **PUT**: replace status of the specified Node

### 14.9.2.1. /api/v1/nodes

Table 14.178. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete collection of Node

**Table 14.179. Query parameters**

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

Parameter	Type	Description
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 14.180. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.181. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Node

**Table 14.182. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.183. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">NodeList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a Node

**Table 14.184. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.185. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Node</b> schema	

**Table 14.186. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Node</b> schema

HTTP code	Response body
201 - Created	<b>Node</b> schema
202 - Accepted	<b>Node</b> schema
401 - Unauthorized	Empty

### 14.9.2.2. /api/v1/watch/nodes

Table 14.187. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method

GET



**Description**

watch individual changes to a list of Node. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.188. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**14.9.2.3. /api/v1/nodes/{name}****Table 14.189. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Node

**Table 14.190. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a Node

**Table 14.191. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.192. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.193. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Node

Table 14.194. HTTP responses

HTTP code	Response body
200 - OK	<b>Node</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Node

**Table 14.195. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.196. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.197. HTTP responses

HTTP code	Response body
200 - OK	<b>Node</b> schema
201 - Created	<b>Node</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Node

Table 14.198. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.199. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Node</b> schema	

Table 14.200. HTTP responses

HTTP code	Response body
200 - OK	<b>Node</b> schema
201 - Created	<b>Node</b> schema
401 - Unauthorized	Empty

#### 14.9.2.4. /api/v1/watch/nodes/{name}

Table 14.201. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Node

Table 14.202. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Node. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.203. HTTP responses**



HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.9.2.5. /api/v1/nodes/{name}/status

Table 14.204. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Node

Table 14.205. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Node

Table 14.206. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">Node</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Node

Table 14.207. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.208. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.209. HTTP responses

HTTP code	Response body
200 - OK	<b>Node</b> schema
201 - Created	<b>Node</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Node

**Table 14.210. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.211. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Node</b> schema	

Table 14.212. HTTP responses

HTTP code	Response body
200 - OK	<b>Node</b> schema
201 - Created	<b>Node</b> schema
401 - Unauthorized	Empty

## 14.10. PERSISTENTVOLUME [V1]

### Description

PersistentVolume (PV) is a storage resource provisioned by an administrator. It is analogous to a node. More info: <https://kubernetes.io/docs/concepts/storage/persistent-volumes>

### Type

**object**

### 14.10.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PersistentVolumeSpec is the specification of a persistent volume.
<b>status</b>	<b>object</b>	PersistentVolumeStatus is the current status of a persistent volume.

#### 14.10.1.1. .spec

##### Description

PersistentVolumeSpec is the specification of a persistent volume.

##### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains all ways the volume can be mounted. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes</a>
<b>awsElasticBlockStore</b>	<b>object</b>	Represents a Persistent Disk resource in AWS.  An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.
<b>azureDisk</b>	<b>object</b>	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

Property	Type	Description
<b>azureFile</b>	<b>object</b>	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
<b>capacity</b>	<b>object (Quantity)</b>	capacity is the description of the persistent volume's resources and capacity. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#capacity">https://kubernetes.io/docs/concepts/storage/persistent-volumes#capacity</a>
<b>cephfs</b>	<b>object</b>	Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.
<b>cinder</b>	<b>object</b>	Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.
<b>claimRef</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.
<b>csi</b>	<b>object</b>	Represents storage that is managed by an external CSI volume driver (Beta feature)
<b>fc</b>	<b>object</b>	Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.
<b>flexVolume</b>	<b>object</b>	FlexPersistentVolumeSource represents a generic persistent volume resource that is provisioned/attached using an exec based plugin.

Property	Type	Description
<b>flocker</b>	<b>object</b>	Represents a Flocker volume mounted by the Flocker agent. One and only one of <code>datasetName</code> and <code>datasetUUID</code> should be set. Flocker volumes do not support ownership management or SELinux relabeling.
<b>gcePersistentDisk</b>	<b>object</b>	Represents a Persistent Disk resource in Google Compute Engine.  A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.
<b>glusterfs</b>	<b>object</b>	Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.
<b>hostPath</b>	<b>object</b>	Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.
<b>iscsi</b>	<b>object</b>	<code>ISCSIPersistentVolumeSource</code> represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.
<b>local</b>	<b>object</b>	Local represents directly-attached storage with node affinity (Beta feature)

Property	Type	Description
<b>mountOptions</b>	<b>array (string)</b>	mountOptions is the list of mount options, e.g. ["ro", "soft"]. Not validated - mount will simply fail if one is invalid. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes/#mount-options">https://kubernetes.io/docs/concepts/storage/persistent-volumes/#mount-options</a>
<b>nfs</b>	<b>object</b>	Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.
<b>nodeAffinity</b>	<b>object</b>	VolumeNodeAffinity defines constraints that limit what nodes this volume can be accessed from.
<b>persistentVolumeReclaimPolicy</b>	<b>string</b>	<p>persistentVolumeReclaimPolicy defines what happens to a persistent volume when released from its claim. Valid options are Retain (default for manually created PersistentVolumes), Delete (default for dynamically provisioned PersistentVolumes), and Recycle (deprecated). Recycle must be supported by the volume plugin underlying this PersistentVolume. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#reclaiming">https://kubernetes.io/docs/concepts/storage/persistent-volumes#reclaiming</a></p> <p>Possible enum values: - <b>"Delete"</b> means the volume will be deleted from Kubernetes on release from its claim. The volume plugin must support Deletion. - <b>"Recycle"</b> means the volume will be recycled back into the pool of unbound persistent volumes on release from its claim. The volume plugin must support Recycling. - <b>"Retain"</b> means the volume will be left in its current phase (Released) for manual reclamation by the administrator. The default policy is Retain.</p>



Property	Type	Description
<b>photonPersistentDisk</b>	<b>object</b>	Represents a Photon Controller persistent disk resource.
<b>portworxVolume</b>	<b>object</b>	PortworxVolumeSource represents a Portworx volume resource.
<b>quobyte</b>	<b>object</b>	Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.
<b>rbd</b>	<b>object</b>	Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.
<b>scaleIO</b>	<b>object</b>	ScaleIOPersistentVolumeSource represents a persistent ScaleIO volume
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of StorageClass to which this persistent volume belongs. Empty value means that this volume does not belong to any StorageClass.
<b>storageos</b>	<b>object</b>	Represents a StorageOS persistent volume resource.
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines if a volume is intended to be used with a formatted filesystem or to remain in raw block state. Value of Filesystem is implied when not included in spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>

Property	Type	Description
<b>vsphereVolume</b>	<b>object</b>	Represents a vSphere volume resource.

### 14.10.1.2. .spec.awsElasticBlockStore

#### Description

Represents a Persistent Disk resource in AWS.

An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

#### Type

**object**

#### Required

- **volumeID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
<b>readOnly</b>	<b>boolean</b>	readOnly value true will force the readOnly setting in VolumeMounts. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>

Property	Type	Description
<b>volumelD</b>	<b>string</b>	volumelD is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>

### 14.10.1.3. .spec.azureDisk

#### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

#### Type

**object**

#### Required

- **diskName**
- **diskURI**

Property	Type	Description
<b>cachingMode</b>	<b>string</b>	cachingMode is the Host Caching mode: None, Read Only, Read Write.  Possible enum values: - <b>"None"</b> - <b>"ReadOnly"</b> - <b>"ReadWrite"</b>
<b>diskName</b>	<b>string</b>	diskName is the Name of the data disk in the blob storage
<b>diskURI</b>	<b>string</b>	diskURI is the URI of data disk in the blob storage
<b>fsType</b>	<b>string</b>	fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

Property	Type	Description
<b>kind</b>	<b>string</b>	kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared  Possible enum values: - <b>"Dedicated"</b> - <b>"Managed"</b> - <b>"Shared"</b>
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

#### 14.10.1.4. .spec.azureFile

##### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **secretName**
- **shareName**

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretName</b>	<b>string</b>	secretName is the name of secret that contains Azure Storage Account Name and Key
<b>secretNamespace</b>	<b>string</b>	secretNamespace is the namespace of the secret that contains Azure Storage Account Name and Key default is the same as the Pod

Property	Type	Description
<b>shareName</b>	<b>string</b>	shareName is the azure Share Name

### 14.10.1.5. .spec.cephfs

#### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

#### Type

**object**

#### Required

- **monitors**

Property	Type	Description
<b>monitors</b>	<b>array (string)</b>	monitors is Required: Monitors is a collection of Ceph monitors More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>path</b>	<b>string</b>	path is Optional: Used as the mounted root, rather than the full Ceph tree, default is /
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretFile</b>	<b>string</b>	secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

Property	Type	Description
<b>user</b>	<b>string</b>	user is Optional: User is the rados user name, default is admin More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

#### 14.10.1.6. .spec.cephfs.secretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.7. .spec.cinder

##### Description

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>secretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>volumeID</b>	<b>string</b>	volumeID used to identify the volume in cinder. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

#### 14.10.1.8. .spec.cinder.secretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.9. .spec.claimRef

## Description

ObjectReference contains enough information to let you inspect or modify the referred object.

## Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>



Property	Type	Description
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

#### 14.10.1.10. .spec.csi

##### Description

Represents storage that is managed by an external CSI volume driver (Beta feature)

##### Type

**object**

##### Required

- **driver**
- **volumeHandle**

Property	Type	Description
<b>controllerExpandSecretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>controllerPublishSecretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume. Required.
<b>fsType</b>	<b>string</b>	fsType to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs".

Property	Type	Description
<b>nodeExpandSecretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>nodePublishSecretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>nodeStageSecretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>readOnly</b>	<b>boolean</b>	readOnly value to pass to ControllerPublishVolumeRequest. Defaults to false (read/write).
<b>volumeAttributes</b>	<b>object (string)</b>	volumeAttributes of the volume to publish.
<b>volumeHandle</b>	<b>string</b>	volumeHandle is the unique volume name returned by the CSI volume plugin's CreateVolume to refer to the volume on all subsequent calls. Required.

#### 14.10.1.11. .spec.csi.controllerExpandSecretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

### 14.10.1.12. .spec.csi.controllerPublishSecretRef

#### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

### 14.10.1.13. .spec.csi.nodeExpandSecretRef

#### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

### 14.10.1.14. .spec.csi.nodePublishSecretRef

#### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.15. .spec.csi.nodeStageSecretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.16. .spec.fc

##### Description

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

Property	Type	Description
<b>lun</b>	<b>integer</b>	lun is Optional: FC target lun number
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>targetWWNs</b>	<b>array (string)</b>	targetWWNs is Optional: FC target worldwide names (WWNs)
<b>wwids</b>	<b>array (string)</b>	wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

#### 14.10.1.17. .spec.flexVolume

##### Description

FlexPersistentVolumeSource represents a generic persistent volume resource that is provisioned/attached using an exec based plugin.

##### Type

**object**

##### Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume.
<b>fsType</b>	<b>string</b>	fsType is the Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
<b>options</b>	<b>object (string)</b>	options is Optional: this field holds extra command options if any.

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

#### 14.10.1.18. .spec.flexVolume.secretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.19. .spec.flocker

##### Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>datasetName</b>	<b>string</b>	datasetName is Name of the dataset stored as metadata → name on the dataset for Flocker should be considered as deprecated

Property	Type	Description
<b>datasetUUID</b>	<b>string</b>	datasetUUID is the UUID of the dataset. This is unique identifier of a Flocker dataset

#### 14.10.1.20. .spec.gcePersistentDisk

##### Description

Represents a Persistent Disk resource in Google Compute Engine.

A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **pdName**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

Property	Type	Description
<b>pdName</b>	<b>string</b>	pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

#### 14.10.1.21. .spec.glusterfs

##### Description

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **endpoints**
- **path**

Property	Type	Description
<b>endpoints</b>	<b>string</b>	endpoints is the endpoint name that details Glusterfs topology. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>



Property	Type	Description
<b>endpointsNamespace</b>	<b>string</b>	endpointsNamespace is the namespace that contains Glusterfs endpoint. If this field is empty, the EndpointNamespace defaults to the same namespace as the bound PVC. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>path</b>	<b>string</b>	path is the Glusterfs volume path. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

#### 14.10.1.22. .spec.hostPath

##### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>

Property	Type	Description
<b>type</b>	<b>string</b>	<p>type for HostPath Volume Defaults to "" More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a></p> <p>Possible enum values: - "" For backwards compatible, leave it empty if unset - <b>"BlockDevice"</b> A block device must exist at the given path - <b>"CharDevice"</b> A character device must exist at the given path - <b>"Directory"</b> A directory must exist at the given path - <b>"DirectoryOrCreate"</b> If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet. - <b>"File"</b> A file must exist at the given path - <b>"FileOrCreate"</b> If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet. - <b>"Socket"</b> A UNIX socket must exist at the given path</p>

### 14.10.1.23. .spec.iscsi

#### Description

ISCSIPersistentVolumeSource represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.

#### Type

**object**

#### Required

- **targetPortal**
- **iqn**
- **lun**

Property	Type	Description
<b>chapAuthDiscovery</b>	<b>boolean</b>	chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication

Property	Type	Description
<b>chapAuthSession</b>	<b>boolean</b>	chapAuthSession defines whether support iSCSI Session CHAP authentication
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#iscsi">https://kubernetes.io/docs/concepts/storage/volumes#iscsi</a>
<b>initiatorName</b>	<b>string</b>	initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <target portal>:<volume name> will be created for the connection.
<b>iqn</b>	<b>string</b>	iqn is Target iSCSI Qualified Name.
<b>iscsiInterface</b>	<b>string</b>	iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
<b>lun</b>	<b>integer</b>	lun is iSCSI Target Lun number.
<b>portals</b>	<b>array (string)</b>	portals is the iSCSI Target Portal List. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
<b>secretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

Property	Type	Description
<b>targetPortal</b>	<b>string</b>	targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

#### 14.10.1.24. .spec.iscsi.secretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.25. .spec.local

##### Description

Local represents directly-attached storage with node affinity (Beta feature)

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. It applies only when the Path is a block device. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default value is to auto-select a filesystem if unspecified.

Property	Type	Description
<b>path</b>	<b>string</b>	path of the full path to the volume on the node. It can be either a directory or block device (disk, partition, ...).

#### 14.10.1.26. .spec.nfs

##### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **server**
- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path that is exported by the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>server</b>	<b>string</b>	server is the hostname or IP address of the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

#### 14.10.1.27. .spec.nodeAffinity

##### Description

VolumeNodeAffinity defines constraints that limit what nodes this volume can be accessed from.

##### Type

**object**

Property	Type	Description
<b>required</b>	<b>object</b>	A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### 14.10.1.28. .spec.nodeAffinity.required

##### Description

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

##### Type

**object**

##### Required

- **nodeSelectorTerms**

Property	Type	Description
<b>nodeSelectorTerms</b>	<b>array</b>	Required. A list of node selector terms. The terms are ORed.
<b>nodeSelectorTerms[]</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

#### 14.10.1.29. .spec.nodeAffinity.required.nodeSelectorTerms

##### Description

Required. A list of node selector terms. The terms are ORed.

##### Type

**array**

#### 14.10.1.30. .spec.nodeAffinity.required.nodeSelectorTerms[]

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.10.1.31. `.spec.nodeAffinity.required.nodeSelectorTerms[].matchExpressions`

##### Description

A list of node selector requirements by node's labels.

##### Type

**array**

#### 14.10.1.32. `.spec.nodeAffinity.required.nodeSelectorTerms[].matchExpressions[]`

##### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

##### Type

**object**

##### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.

Property	Type	Description
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 14.10.1.33. .spec.nodeAffinity.required.nodeSelectorTerms[].matchFields

##### Description

A list of node selector requirements by node's fields.

##### Type

**array**

#### 14.10.1.34. .spec.nodeAffinity.required.nodeSelectorTerms[].matchFields[]

##### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

##### Type

**object**

##### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.



Property	Type	Description
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 14.10.1.35. .spec.photonPersistentDisk

##### Description

Represents a Photon Controller persistent disk resource.

##### Type

**object**

##### Required

- **pdID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>pdID</b>	<b>string</b>	pdID is the ID that identifies Photon Controller persistent disk

#### 14.10.1.36. .spec.portworxVolume

##### Description

PortworxVolumeSource represents a Portworx volume resource.

#### Type

**object**

#### Required

- **volumelD**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>volumelD</b>	<b>string</b>	volumelD uniquely identifies a Portworx volume

#### 14.10.1.37. .spec.quobyte

#### Description

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

#### Type

**object**

#### Required

- **registry**
- **volume**

Property	Type	Description
<b>group</b>	<b>string</b>	group to map volume access to. Default is no group
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.

Property	Type	Description
<b>registry</b>	<b>string</b>	registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
<b>tenant</b>	<b>string</b>	tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
<b>user</b>	<b>string</b>	user to map volume access to Defaults to serviceaccount user
<b>volume</b>	<b>string</b>	volume is a string that references an already created Quobyte volume by name.

#### 14.10.1.38. .spec.rbd

##### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**
- **image**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#rbd">https://kubernetes.io/docs/concepts/storage/volumes#rbd</a>

Property	Type	Description
<b>image</b>	<b>string</b>	image is the rados image name. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>keyring</b>	<b>string</b>	keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>monitors</b>	<b>array (string)</b>	monitors is a collection of Ceph monitors. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>pool</b>	<b>string</b>	pool is the rados pool name. Default is rbd. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace
<b>user</b>	<b>string</b>	user is the rados user name. Default is admin. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

#### 14.10.1.39. .spec.rbd.secretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.40. .spec.scaleIO

##### Description

ScaleIOPersistentVolumeSource represents a persistent ScaleIO volume

##### Type

**object**

##### Required

- **gateway**
- **system**
- **secretRef**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs"
<b>gateway</b>	<b>string</b>	gateway is the host address of the ScaleIO API Gateway.
<b>protectionDomain</b>	<b>string</b>	protectionDomain is the name of the ScaleIO Protection Domain for the configured storage.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

Property	Type	Description
<b>sslEnabled</b>	<b>boolean</b>	sslEnabled is the flag to enable/disable SSL communication with Gateway, default false
<b>storageMode</b>	<b>string</b>	storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
<b>storagePool</b>	<b>string</b>	storagePool is the ScaleIO Storage Pool associated with the protection domain.
<b>system</b>	<b>string</b>	system is the name of the storage system as configured in ScaleIO.
<b>volumeName</b>	<b>string</b>	volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source.

#### 14.10.1.41. .spec.scaleIO.secretRef

##### Description

SecretReference represents a Secret Reference. It has enough information to retrieve secret in any namespace

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is unique within a namespace to reference a secret resource.
<b>namespace</b>	<b>string</b>	namespace defines the space within which the secret name must be unique.

#### 14.10.1.42. .spec.storageeos

##### Description

Represents a StorageOS persistent volume resource.

##### Type

**object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.
<b>volumeName</b>	<b>string</b>	volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
<b>volumeNamespace</b>	<b>string</b>	volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

**14.10.1.43. .spec.storageos.secretRef****Description**

ObjectReference contains enough information to let you inspect or modify the referred object.

**Type**

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>



Property	Type	Description
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

#### 14.10.1.44. .spec.vsphereVolume

##### Description

Represents a vSphere volume resource.

##### Type

**object**

##### Required

- **volumePath**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>storagePolicyID</b>	<b>string</b>	storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
<b>storagePolicyName</b>	<b>string</b>	storagePolicyName is the storage Policy Based Management (SPBM) profile name.
<b>volumePath</b>	<b>string</b>	volumePath is the path that identifies vSphere volume vmdk

#### 14.10.1.45. .status

##### Description

PersistentVolumeStatus is the current status of a persistent volume.

##### Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	message is a human-readable message indicating details about why the volume is in this state.
<b>phase</b>	<b>string</b>	<p>phase indicates if a volume is available, bound to a claim, or released by a claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#phase">https://kubernetes.io/docs/concepts/storage/persistent-volumes#phase</a></p> <p>Possible enum values: -  <b>"Available"</b> used for PersistentVolumes that are not yet bound Available volumes are held by the binder and matched to PersistentVolumeClaims -  <b>"Bound"</b> used for PersistentVolumes that are bound - <b>"Failed"</b> used for PersistentVolumes that failed to be correctly recycled or deleted after being released from a claim - <b>"Pending"</b> used for PersistentVolumes that are not available - <b>"Released"</b> used for PersistentVolumes where the bound PersistentVolumeClaim was deleted released volumes must be recycled before becoming available again this phase is used by the persistent volume claim binder to signal to another process to reclaim the resource</p>
<b>reason</b>	<b>string</b>	reason is a brief CamelCase string that describes any failure and is meant for machine parsing and tidy display in the CLI.

### 14.10.2. API endpoints

The following API endpoints are available:

- **/api/v1/persistentvolumes**
  - **DELETE:** delete collection of PersistentVolume
  - **GET:** list or watch objects of kind PersistentVolume

- **POST**: create a PersistentVolume
- **/api/v1/watch/persistentvolumes**
  - **GET**: watch individual changes to a list of PersistentVolume. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/persistentvolumes/{name}**
  - **DELETE**: delete a PersistentVolume
  - **GET**: read the specified PersistentVolume
  - **PATCH**: partially update the specified PersistentVolume
  - **PUT**: replace the specified PersistentVolume
- **/api/v1/watch/persistentvolumes/{name}**
  - **GET**: watch changes to an object of kind PersistentVolume. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/persistentvolumes/{name}/status**
  - **GET**: read status of the specified PersistentVolume
  - **PATCH**: partially update status of the specified PersistentVolume
  - **PUT**: replace status of the specified PersistentVolume

#### 14.10.2.1. /api/v1/persistentvolumes

Table 14.213. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### **DELETE**

#### Description

delete collection of PersistentVolume

Table 14.214. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>

Parameter	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 14.215. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.216. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind PersistentVolume

Table 14.217. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.218. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PersistentVolumeList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a PersistentVolume

**Table 14.219. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.220. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PersistentVolume</b> schema	

**Table 14.221. HTTP responses**

HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema

HTTP code	Response body
201 - Created	<b>PersistentVolume</b> schema
202 - Accepted	<b>PersistentVolume</b> schema
401 - Unauthorized	Empty

### 14.10.2.2. /api/v1/watch/persistentvolumes

Table 14.222. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PersistentVolume. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.223. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.10.2.3. /api/v1/persistentvolumes/{name}

Table 14.224. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PersistentVolume

Table 14.225. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a PersistentVolume

Table 14.226. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.227. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.228. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema
202 - Accepted	<b>PersistentVolume</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified PersistentVolume

Table 14.229. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema



HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified PersistentVolume

**Table 14.230. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.231. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.232. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema
201 - Created	<b>PersistentVolume</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified PersistentVolume

Table 14.233. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.234. Body parameters

Parameter	Type	Description
<b>body</b>	<b>PersistentVolume</b> schema	

Table 14.235. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema
201 - Created	<b>PersistentVolume</b> schema
401 - Unauthorized	Empty

#### 14.10.2.4. /api/v1/watch/persistentvolumes/{name}

Table 14.236. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PersistentVolume

Table 14.237. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind PersistentVolume. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 14.238. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.10.2.5. /api/v1/persistentvolumes/{name}/status

Table 14.239. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PersistentVolume

Table 14.240. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified PersistentVolume

Table 14.241. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">PersistentVolume</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified PersistentVolume

Table 14.242. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.243. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.244. HTTP responses



HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema
201 - Created	<b>PersistentVolume</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified PersistentVolume

**Table 14.245. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.246. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PersistentVolume</b> schema	

Table 14.247. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolume</b> schema
201 - Created	<b>PersistentVolume</b> schema
401 - Unauthorized	Empty

## 14.11. PERSISTENTVOLUMECLAIM [V1]

### Description

PersistentVolumeClaim is a user's request for and claim to a persistent volume

### Type

**object**

### 14.11.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes
<b>status</b>	<b>object</b>	PersistentVolumeClaimStatus is the current status of a persistent volume claim.

#### 14.11.1.1. .spec

##### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

##### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the desired access modes the volume should have. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>
<b>dataSource</b>	<b>object</b>	TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

Property	Type	Description
<b>dataSourceRef</b>	<b>object</b>	<p>dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: * While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. * While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. * While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.</p>

Property	Type	Description
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over volumes to consider for binding.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of the StorageClass required by the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1</a>
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>
<b>volumeName</b>	<b>string</b>	volumeName is the binding reference to the PersistentVolume backing this claim.

#### 14.11.1.2. .spec.dataSource

##### Description

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

##### Type

**object**

##### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

### 14.11.1.3. .spec.dataSourceRef

#### Description

`dataSourceRef` specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a `PersistentVolumeClaim` object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the `dataSource` field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in `dataSourceRef`, both fields (`dataSource` and `dataSourceRef`) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in `dataSourceRef`, `dataSource` isn't set to the same value and must be empty. There are three important differences between `dataSource` and `dataSourceRef`: \* While `dataSource` only allows two specific types of objects, `dataSourceRef` allows any non-core object, as well as `PersistentVolumeClaim` objects. \* While `dataSource` ignores disallowed values (dropping them), `dataSourceRef` preserves all values, and generates an error if a disallowed value is specified. \* While `dataSource` only allows local objects, `dataSourceRef` allows objects in any namespaces. (Beta) Using this field requires the `AnyVolumeDataSource` feature gate to be enabled. (Alpha) Using the namespace field of `dataSourceRef` requires the `CrossNamespaceVolumeDataSource` feature gate to be enabled.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced
<b>namespace</b>	<b>string</b>	Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

#### 14.11.1.4. .spec.resources

##### Description

ResourceRequirements describes the compute resource requirements.

##### Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>

Property	Type	Description
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.11.1.5. .spec.resources.claims

##### Description

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.11.1.6. .spec.resources.claims[]

##### Description

ResourceClaim references one entry in PodSpec.ResourceClaims.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.11.1.7. .status

##### Description

`PersistentVolumeClaimStatus` is the current status of a persistent volume claim.

##### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	<code>accessModes</code> contains the actual access modes the volume backing the PVC has. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>
<b>allocatedResources</b>	<b>object (Quantity)</b>	<code>allocatedResources</code> is the storage resource within <code>AllocatedResources</code> tracks the capacity allocated to a PVC. It may be larger than the actual capacity when a volume expansion operation is requested. For storage quota, the larger value from <code>allocatedResources</code> and <code>PVC.spec.resources</code> is used. If <code>allocatedResources</code> is not set, <code>PVC.spec.resources</code> alone is used for quota calculation. If a volume expansion capacity request is lowered, <code>allocatedResources</code> is only lowered if there are no expansion operations in progress and if the actual volume capacity is equal or lower than the requested capacity. This is an alpha field and requires enabling <code>RecoverVolumeExpansionFailure</code> feature.

Property	Type	Description
<b>capacity</b>	<b>object (Quantity)</b>	capacity represents the actual resources of the underlying volume.
<b>conditions</b>	<b>array</b>	conditions is the current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
<b>conditions[]</b>	<b>object</b>	PersistentVolumeClaimCondition contains details about state of pvc
<b>phase</b>	<b>string</b>	phase represents the current phase of PersistentVolumeClaim.  Possible enum values: - <b>"Bound"</b> used for PersistentVolumeClaims that are bound - <b>"Lost"</b> used for PersistentVolumeClaims that lost their underlying PersistentVolume. The claim was bound to a PersistentVolume and this volume does not exist any longer and all data on it was lost. - <b>"Pending"</b> used for PersistentVolumeClaims that are not yet bound

Property	Type	Description
<b>resizeStatus</b>	<b>string</b>	<p>resizeStatus stores status of resize operation. ResizeStatus is not set by default but when expansion is complete resizeStatus is set to empty string by resize controller or kubelet. This is an alpha field and requires enabling RecoverVolumeExpansionFailure feature.</p> <p>Possible enum values: - "" When expansion is complete, the empty string is set by resize controller or kubelet. - <b>"ControllerExpansionFailed"</b> State set when expansion has failed in resize controller with a terminal error. Transient errors such as timeout should not set this status and should leave ResizeStatus unmodified, so as resize controller can resume the volume expansion. - <b>"ControllerExpansionInProgress"</b> State set when resize controller starts expanding the volume in control-plane - <b>"NodeExpansionFailed"</b> State set when expansion has failed in kubelet with a terminal error. Transient errors don't set NodeExpansionFailed. - <b>"NodeExpansionInProgress"</b> State set when kubelet starts expanding the volume. - <b>"NodeExpansionPending"</b> State set when resize controller has finished expanding the volume but further expansion is needed on the node.</p>

#### 14.11.1.8. .status.conditions

##### Description

conditions is the current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.

##### Type

**array**

### 14.11.1.9. .status.conditions[]

#### Description

PersistentVolumeClaimCondition contains details about state of pvc

#### Type

**object**

#### Required

- **type**
- **status**

Property	Type	Description
<b>lastProbeTime</b>	<b>Time</b>	lastProbeTime is the time we probed the condition.
<b>lastTransitionTime</b>	<b>Time</b>	lastTransitionTime is the time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	message is the human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	reason is a unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
<b>status</b>	<b>string</b>	
<b>type</b>	<b>string</b>	

### 14.11.2. API endpoints

The following API endpoints are available:

- **/api/v1/persistentvolumeclaims**
  - **GET**: list or watch objects of kind PersistentVolumeClaim
- **/api/v1/watch/persistentvolumeclaims**
  - **GET**: watch individual changes to a list of PersistentVolumeClaim. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/persistentvolumeclaims**

- **DELETE:** delete collection of PersistentVolumeClaim
- **GET:** list or watch objects of kind PersistentVolumeClaim
- **POST:** create a PersistentVolumeClaim
- **/api/v1/watch/namespaces/{namespace}/persistentvolumeclaims**
  - **GET:** watch individual changes to a list of PersistentVolumeClaim. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}**
  - **DELETE:** delete a PersistentVolumeClaim
  - **GET:** read the specified PersistentVolumeClaim
  - **PATCH:** partially update the specified PersistentVolumeClaim
  - **PUT:** replace the specified PersistentVolumeClaim
- **/api/v1/watch/namespaces/{namespace}/persistentvolumeclaims/{name}**
  - **GET:** watch changes to an object of kind PersistentVolumeClaim. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}/status**
  - **GET:** read status of the specified PersistentVolumeClaim
  - **PATCH:** partially update status of the specified PersistentVolumeClaim
  - **PUT:** replace status of the specified PersistentVolumeClaim

#### 14.11.2.1. /api/v1/persistentvolumeclaims

Table 14.248. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind PersistentVolumeClaim

**Table 14.249. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">PersistentVolumeClaimList</a> schema



HTTP code	Response body
401 - Unauthorized	Empty

### 14.11.2.2. /api/v1/watch/persistentvolumeclaims

Table 14.250. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PersistentVolumeClaim. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.251. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.11.2.3. /api/v1/namespaces/{namespace}/persistentvolumeclaims

Table 14.252. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.253. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of PersistentVolumeClaim

Table 14.254. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.255. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.256. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind PersistentVolumeClaim

Table 14.257. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.258. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PersistentVolumeClaimList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a PersistentVolumeClaim

**Table 14.259. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.260. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PersistentVolumeClaim</b> schema	

**Table 14.261. HTTP responses**

HTTP code	Response body
200 - OK	<b>PersistentVolumeClaim</b> schema

HTTP code	Response body
201 - Created	<a href="#">PersistentVolumeClaim</a> schema
202 - Accepted	<a href="#">PersistentVolumeClaim</a> schema
401 - Unauthorized	Empty

#### 14.11.2.4. /api/v1/watch/namespaces/{namespace}/persistentvolumeclaims

Table 14.262. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.263. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PersistentVolumeClaim. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.264. HTTP responses



HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.11.2.5. /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}

Table 14.265. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PersistentVolumeClaim
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.266. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a PersistentVolumeClaim

Table 14.267. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.268. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.269. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolumeClaim</b> schema
202 - Accepted	<b>PersistentVolumeClaim</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified PersistentVolumeClaim

Table 14.270. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolumeClaim</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified PersistentVolumeClaim

**Table 14.271. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.272. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.273. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolumeClaim</b> schema
201 - Created	<b>PersistentVolumeClaim</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified PersistentVolumeClaim

Table 14.274. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.275. Body parameters

Parameter	Type	Description
<b>body</b>	<b>PersistentVolumeClaim</b> schema	

Table 14.276. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolumeClaim</b> schema
201 - Created	<b>PersistentVolumeClaim</b> schema
401 - Unauthorized	Empty

#### 14.11.2.6. /api/v1/watch/namespaces/{namespace}/persistentvolumeclaims/{name}

Table 14.277. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PersistentVolumeClaim
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 14.278. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind PersistentVolumeClaim. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.279. HTTP responses**



HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.11.2.7. /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}/status

Table 14.280. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PersistentVolumeClaim
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.281. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified PersistentVolumeClaim

Table 14.282. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">PersistentVolumeClaim</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified PersistentVolumeClaim

Table 14.283. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.284. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.285. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PersistentVolumeClaim</a> schema
201 - Created	<a href="#">PersistentVolumeClaim</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified PersistentVolumeClaim

**Table 14.286. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.287. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PersistentVolumeClaim</b> schema	

Table 14.288. HTTP responses

HTTP code	Response body
200 - OK	<b>PersistentVolumeClaim</b> schema
201 - Created	<b>PersistentVolumeClaim</b> schema
401 - Unauthorized	Empty

## 14.12. POD [V1]

### Description

Pod is a collection of containers that can run on a host. This resource is created by clients and scheduled onto hosts.

### Type

**object**

### 14.12.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PodSpec is a description of a pod.
<b>status</b>	<b>object</b>	PodStatus represents information about the status of a pod. Status may trail the actual state of a system, especially if the node that hosts the pod cannot contact the control plane.

#### 14.12.1.1. .spec

##### Description

PodSpec is a description of a pod.

##### Type

**object**

##### Required

- **containers**

Property	Type	Description
<b>activeDeadlineSeconds</b>	<b>integer</b>	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.

Property	Type	Description
<b>affinity</b>	<b>object</b>	Affinity is a group of affinity scheduling rules.
<b>automountServiceAccountToken</b>	<b>boolean</b>	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
<b>containers</b>	<b>array</b>	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
<b>containers[]</b>	<b>object</b>	A single application container that you want to run within a pod.
<b>dnsConfig</b>	<b>object</b>	PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

Property	Type	Description
<b>dnsPolicy</b>	<b>string</b>	<p>Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.</p> <p>Possible enum values: - <b>"ClusterFirst"</b> indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"ClusterFirstWithHostNet"</b> indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"Default"</b> indicates that the pod should use the default (as determined by kubelet) DNS settings. - <b>"None"</b> indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig.</p>
<b>enableServiceLinks</b>	<b>boolean</b>	<p>EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.</p>

Property	Type	Description
<b>ephemeralContainers</b>	<b>array</b>	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's <code>ephemeralcontainers</code> subresource.
<b>ephemeralContainers[]</b>	<b>object</b>	<p>An <code>EphemeralContainer</code> is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation.</p> <p>To add an ephemeral container, use the <code>ephemeralcontainers</code> subresource of an existing Pod. Ephemeral containers may not be removed or restarted.</p>
<b>hostAliases</b>	<b>array</b>	<code>HostAliases</code> is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non- <code>hostNetwork</code> pods.
<b>hostAliases[]</b>	<b>object</b>	<code>HostAlias</code> holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.
<b>hostIPC</b>	<b>boolean</b>	Use the host's <code>ipc</code> namespace. Optional: Default to false.



Property	Type	Description
<b>hostNetwork</b>	<b>boolean</b>	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
<b>hostPID</b>	<b>boolean</b>	Use the host's pid namespace. Optional: Default to false.
<b>hostUsers</b>	<b>boolean</b>	Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new users is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature.
<b>hostname</b>	<b>string</b>	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
<b>imagePullSecrets</b>	<b>array</b>	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod">https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod</a>
<b>imagePullSecrets[]</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

Property	Type	Description
<b>initContainers</b>	<b>array</b>	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/init-containers/">https://kubernetes.io/docs/concepts/workloads/pods/init-containers/</a>
<b>initContainers[]</b>	<b>object</b>	A single application container that you want to run within a pod.
<b>nodeName</b>	<b>string</b>	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
<b>nodeSelector</b>	<b>object (string)</b>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a>

Property	Type	Description
<b>os</b>	<b>object</b>	PodOS defines the OS parameters of a pod.
<b>overhead</b>	<b>object (Quantity)</b>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md">https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md</a>
<b>preemptionPolicy</b>	<b>string</b>	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.  Possible enum values: - <b>"Never"</b> means that pod never preempts other pods with lower priority. - <b>"PreemptLowerPriority"</b> means that pod can preempt other pods with lower priority.
<b>priority</b>	<b>integer</b>	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

Property	Type	Description
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
<b>readinessGates</b>	<b>array</b>	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: <a href="https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates">https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates</a>
<b>readinessGates[]</b>	<b>object</b>	PodReadinessGate contains the reference to a pod condition
<b>resourceClaims</b>	<b>array</b>	ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable.

Property	Type	Description
<b>resourceClaims[]</b>	<b>object</b>	PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.
<b>restartPolicy</b>	<b>string</b>	Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy</a>  Possible enum values: - <b>"Always"</b> - <b>"Never"</b> - <b>"OnFailure"</b>
<b>runtimeClassName</b>	<b>string</b>	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class">https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class</a>
<b>schedulerName</b>	<b>string</b>	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.

Property	Type	Description
<b>schedulingGates</b>	<b>array</b>	<p>SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.</p> <p>SchedulingGates can only be set at pod creation time, and be removed only afterwards.</p> <p>This is a beta feature enabled by the PodSchedulingReadiness feature gate.</p>
<b>schedulingGates[]</b>	<b>object</b>	PodSchedulingGate is associated to a Pod to guard its scheduling.
<b>securityContext</b>	<b>object</b>	PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.
<b>serviceAccount</b>	<b>string</b>	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
<b>serviceAccountName</b>	<b>string</b>	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: <a href="https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/">https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/</a>

Property	Type	Description
<b>setHostnameAsFQDN</b>	<b>boolean</b>	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
<b>shareProcessNamespace</b>	<b>boolean</b>	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
<b>subdomain</b>	<b>string</b>	If specified, the fully qualified Pod hostname will be "<hostname>.<subdomain>.<pod namespace>.svc.<cluster domain>". If not specified, the pod will not have a domainname at all.

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
<b>tolerations</b>	<b>array</b>	If specified, the pod's tolerations.
<b>tolerations[]</b>	<b>object</b>	The pod this Tolerant is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.
<b>topologySpreadConstraints</b>	<b>array</b>	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
<b>topologySpreadConstraints[]</b>	<b>object</b>	TopologySpreadConstraint specifies how to spread matching pods among the given topology.
<b>volumes</b>	<b>array</b>	List of volumes that can be mounted by containers belonging to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes">https://kubernetes.io/docs/concepts/storage/volumes</a>
<b>volumes[]</b>	<b>object</b>	Volume represents a named volume in a pod that may be accessed by any container in the pod.



### 14.12.1.2. .spec.affinity

#### Description

Affinity is a group of affinity scheduling rules.

#### Type

**object**

Property	Type	Description
<b>nodeAffinity</b>	<b>object</b>	Node affinity is a group of node affinity scheduling rules.
<b>podAffinity</b>	<b>object</b>	Pod affinity is a group of inter pod affinity scheduling rules.
<b>podAntiAffinity</b>	<b>object</b>	Pod anti affinity is a group of inter pod anti affinity scheduling rules.

### 14.12.1.3. .spec.affinity.nodeAffinity

#### Description

Node affinity is a group of node affinity scheduling rules.

#### Type

**object**

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>object</b>	A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### 14.12.1.4. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution`

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, `requiredDuringScheduling` affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding `matchExpressions`; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.12.1.5. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[]`

##### Description

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

##### Type

**object**

##### Required

- **weight**
- **preference**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>preference</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.
<b>weight</b>	<b>integer</b>	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

#### 14.12.1.6. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].pre`

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.12.1.7. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].pre`

##### Description

A list of node selector requirements by node's labels.

##### Type

**array**

### 14.12.1.8. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].pre`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.12.1.9. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].pre`

#### Description

A list of node selector requirements by node's fields.

#### Type

**array**

### 14.12.1.10. `.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution[].pr`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 14.12.1.11. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution

##### Description

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### Type

**object**

#### Required

- **nodeSelectorTerms**

Property	Type	Description
<b>nodeSelectorTerms</b>	<b>array</b>	Required. A list of node selector terms. The terms are ORed.
<b>nodeSelectorTerms[]</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

#### 14.12.1.12. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms

##### Description

Required. A list of node selector terms. The terms are ORed.

##### Type

**array**

#### 14.12.1.13. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.12.1.14. .spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.nodeSelectorTerms[]

##### Description

A list of node selector requirements by node's labels.

#### Type

**array**

### 14.12.1.15. `.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.node`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.12.1.16. `.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.node`

#### Description

A list of node selector requirements by node's fields.

#### Type

**array**

### 14.12.1.17. `.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.node`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.12.1.18. `.spec.affinity.podAffinity`

#### Description

Pod affinity is a group of inter pod affinity scheduling rules.

#### Type

**object**



Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

#### 14.12.1.19. .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.12.1.20. .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

##### Type

**object**

##### Required

- **weight**
- **podAffinityTerm**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>podAffinityTerm</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

#### 14.12.1.21. .spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution[].pod

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} matches all namespaces.

Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.12.1.22. `.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution`

##### Description

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

##### Type

**array**

#### 14.12.1.23. `.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecution[]`

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.12.1.24. .spec.affinity.podAntiAffinity

##### Description

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

#### 14.12.1.25. .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.12.1.26. .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

##### Type

**object**

##### Required

- **weight**
- **podAffinityTerm**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>podAffinityTerm</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

#### 14.12.1.27. .spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringExecution[]

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} matches all namespaces.



Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.12.1.28. `.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution`

##### Description

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each `podAffinityTerm` are intersected, i.e. all terms must be satisfied.

##### Type

**array**

#### 14.12.1.29. `.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringExecution[]`

##### Description

Defines a set of pods (namely those matching the `labelSelector` relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key `<topologyKey>` matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.12.1.30. .spec.containers

##### Description

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

##### Type

**array**

#### 14.12.1.31. .spec.containers[]

##### Description

A single application container that you want to run within a pod.

Type

**object**

Required

- **name**

Property	Type	Description
<b>args</b>	<b>array (string)</b>	Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references <code>\$(VAR_NAME)</code> are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>\$</code> are reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. " <code>\$(VAR_NAME)</code> " will produce the string literal " <code>\$(VAR_NAME)</code> ". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>
<b>env</b>	<b>array</b>	List of environment variables to set in the container. Cannot be updated.
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps

Property	Type	Description
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
<b>imagePullPolicy</b>	<b>string</b>	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a>  Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present
<b>lifecycle</b>	<b>object</b>	Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.
<b>livenessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
<b>ports</b>	<b>array</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

Property	Type	Description
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.



Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.12.1.32. .spec.containers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.12.1.33. .spec.containers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.12.1.34. `.spec.containers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources ( <code>cpu</code> , <code>memory</code> ) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.12.1.35. `.spec.containers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a `ConfigMap`.

Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.12.1.36. `.spec.containers[].env[].valueFrom.fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.12.1.37. `.spec.containers[].env[].valueFrom.resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.12.1.38. .spec.containers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.12.1.39. .spec.containers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.12.1.40. .spec.containers[].envFrom[]

##### Description

EnvFromSource represents the source of a set of ConfigMaps

### Type

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>object</b>	SecretEnvSource selects a Secret to populate the environment variables with.  The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### 14.12.1.41. .spec.containers[].envFrom[].configMapRef

### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 14.12.1.42. .spec.containers[].envFrom[].secretRef

## Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.12.1.43. .spec.containers[].lifecycle

## Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

## Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.12.1.44. .spec.containers[].lifecycle.postStart

## Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

**14.12.1.45. .spec.containers[].lifecycle.postStart.exec****Description**

ExecAction describes a "run in container" action.

**Type**

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**14.12.1.46. .spec.containers[].lifecycle.postStart.httpGet****Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

**object**

**Required**

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.12.1.47. .spec.containers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.12.1.48. .spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**



Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.49. `..spec.containers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.50. `..spec.containers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.12.1.51. .spec.containers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.52. .spec.containers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.12.1.53. .spec.containers[].lifecycle.preStop.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.12.1.54. .spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.55. .spec.containers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.56. .spec.containers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.57. .spec.containers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.58. .spec.containers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.59. .spec.containers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.12.1.60. .spec.containers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.12.1.61. .spec.containers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required



- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.62. `.spec.containers[].livenessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.63. `.spec.containers[].ports`

##### Description

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <https://github.com/kubernetes/kubernetes/issues/108255>. Cannot be updated.

##### Type

**array**

#### 14.12.1.64. `.spec.containers[].ports[]`

**Description**

ContainerPort represents a network port in a single container.

**Type**

**object**

**Required**

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

**14.12.1.65. .spec.containers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.66. `.spec.containers[].readinessProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.67. .spec.containers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.68. .spec.containers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.12.1.69. `..spec.containers[].readinessProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.12.1.70. `..spec.containers[].readinessProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.71. `.spec.containers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.72. `.spec.containers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.12.1.73. `.spec.containers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.12.1.74. .spec.containers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>



Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.75. `.spec.containers[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.76. `.spec.containers[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.77. `.spec.containers[].securityContext`

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.12.1.78. .spec.containers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities

Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.12.1.79. `.spec.containers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.12.1.80. `.spec.containers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.12.1.81. .spec.containers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.12.1.82. .spec.containers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket



Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.83. .spec.containers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.84. .spec.containers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.85. .spec.containers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.12.1.86. `.spec.containers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.12.1.87. `.spec.containers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.88. `.spec.containers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.89. `.spec.containers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.12.1.90. `.spec.containers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.12.1.91. .spec.containers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Cannot be updated.

## Type

**array**

**14.12.1.92. .spec.containers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '!'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). -  <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

### 14.12.1.93. .spec.dnsConfig

#### Description

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

#### Type

**object**

Property	Type	Description
<b>nameservers</b>	<b>array (string)</b>	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
<b>options</b>	<b>array</b>	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
<b>options[]</b>	<b>object</b>	PodDNSConfigOption defines DNS resolver options of a pod.
<b>searches</b>	<b>array (string)</b>	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

### 14.12.1.94. .spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

**array**

**14.12.1.95. .spec.dnsConfig.options[]****Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Required.
<b>value</b>	<b>string</b>	

**14.12.1.96. .spec.ephemeralContainers****Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

**array**

**14.12.1.97. .spec.ephemeralContainers[]****Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

**object**

**Required**

- **name**



Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Endpoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array</b>	List of environment variables to set in the container. Cannot be updated.

Property	Type	Description
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a>
<b>imagePullPolicy</b>	<b>string</b>	<p>Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a></p> <p>Possible enum values: -  <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail If the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present</p>

Property	Type	Description
<b>lifecycle</b>	<b>object</b>	Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.
<b>livenessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>name</b>	<b>string</b>	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
<b>ports</b>	<b>array</b>	Ports are not allowed for ephemeral containers.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.

Property	Type	Description
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

Property	Type	Description
<b>targetContainerName</b>	<b>string</b>	<p>If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.</p> <p>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined.</p>
<b>terminationMessagePath</b>	<b>string</b>	<p>Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to <code>/dev/termination-log</code>. Cannot be updated.</p>

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.12.1.98. .spec.ephemeralContainers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.12.1.99. .spec.ephemeralContainers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.12.1.100. `.spec.ephemeralContainers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources ( <code>cpu</code> , <code>memory</code> ) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.12.1.101. `.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a `ConfigMap`.



Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.12.1.102. `.spec.ephemeralContainers[].env[].valueFrom.fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.12.1.103. `.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.12.1.104. .spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.12.1.105. .spec.ephemeralContainers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.12.1.106. .spec.ephemeralContainers[].envFrom[]

##### Description

EnvFromSource represents the source of a set of ConfigMaps

#### Type

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>object</b>	SecretEnvSource selects a Secret to populate the environment variables with.  The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### 14.12.1.107. .spec.ephemeralContainers[].envFrom[].configMapRef

##### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  
The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 14.12.1.108. .spec.ephemeralContainers[].envFrom[].secretRef

## Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

## 14.12.1.109. .spec.ephemeralContainers[].lifecycle

### Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

### Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## 14.12.1.110. .spec.ephemeralContainers[].lifecycle.postStart

### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

**14.12.1.111. .spec.ephemeralContainers[].lifecycle.postStart.exec****Description**

ExecAction describes a "run in container" action.

**Type**

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**14.12.1.112. .spec.ephemeralContainers[].lifecycle.postStart.httpGet****Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

**object**

**Required**

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.12.1.113. .spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.12.1.114. .spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.115. `..spec.ephemeralContainers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.116. `..spec.ephemeralContainers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.12.1.117. .spec.ephemeralContainers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.118. .spec.ephemeralContainers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.



Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.12.1.119. .spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.12.1.120. .spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.121. .spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.122. .spec.ephemeralContainers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.123. .spec.ephemeralContainers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.124. .spec.ephemeralContainers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.125. .spec.ephemeralContainers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.12.1.126. .spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.12.1.127. .spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.128. .spec.ephemeralContainers[].livenessProbe.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.129. .spec.ephemeralContainers[].ports

##### Description

Ports are not allowed for ephemeral containers.

##### Type

**array**

#### 14.12.1.130. .spec.ephemeralContainers[].ports[]

##### Description

ContainerPort represents a network port in a single container.

##### Type

**object**

## Required

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.

**14.12.1.131. .spec.ephemeralContainers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type****object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.



Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.132. .spec.ephemeralContainers[].readinessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.133. .spec.ephemeralContainers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.134. .spec.ephemeralContainers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.12.1.135. .spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.12.1.136. .spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.137. .spec.ephemeralContainers[].readinessProbe.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.138. .spec.ephemeralContainers[].resizePolicy

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.12.1.139. .spec.ephemeralContainers[].resizePolicy[]

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.12.1.140. .spec.ephemeralContainers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.141. `.spec.ephemeralContainers[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.142. `.spec.ephemeralContainers[].resources.claims[]`

##### Description

`ResourceClaim` references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.143. `.spec.ephemeralContainers[].securityContext`

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.



Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.12.1.144. .spec.ephemeralContainers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities

Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.12.1.145. `.spec.ephemeralContainers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.12.1.146. `.spec.ephemeralContainers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.12.1.147. .spec.ephemeralContainers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.12.1.148. .spec.ephemeralContainers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.149. .spec.ephemeralContainers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.150. .spec.ephemeralContainers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.151. .spec.ephemeralContainers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**



## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.12.1.152. `.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.12.1.153. `.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.154. `.spec.ephemeralContainers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.155. `.spec.ephemeralContainers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.12.1.156. `.spec.ephemeralContainers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.12.1.157. .spec.ephemeralContainers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

## Type

**array**

**14.12.1.158. .spec.ephemeralContainers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '.'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). -  <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.12.1.159. .spec.hostAliases

##### Description

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.

##### Type

**array**

#### 14.12.1.160. .spec.hostAliases[]

##### Description

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

##### Type

**object**

Property	Type	Description
<b>hostnames</b>	<b>array (string)</b>	Hostnames for the above IP address.
<b>ip</b>	<b>string</b>	IP address of the host file entry.

#### 14.12.1.161. .spec.imagePullSecrets

##### Description

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info:

<https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod>

##### Type

**array**

#### 14.12.1.162. .spec.imagePullSecrets[]

## Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

### 14.12.1.163. .spec.initContainers

## Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <https://kubernetes.io/docs/concepts/workloads/pods/init-containers/>

## Type

**array**

### 14.12.1.164. .spec.initContainers[]

## Description

A single application container that you want to run within a pod.

## Type

**object**

## Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The container image's CMD is used if this is not provided. Variable references <code>\$(VAR_NAME)</code> are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>\$\$</code> are reduced to a single <code>\$</code>, which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. "<code>\$(VAR_NAME)</code>" will produce the string literal <code>\$(VAR_NAME)</code>. Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Endpoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references <code>\$(VAR_NAME)</code> are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>\$\$</code> are reduced to a single <code>\$</code>, which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. "<code>\$(VAR_NAME)</code>" will produce the string literal <code>\$(VAR_NAME)</code>. Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array</b>	List of environment variables to set in the container. Cannot be updated.

Property	Type	Description
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.



Property	Type	Description
<b>imagePullPolicy</b>	<b>string</b>	<p>Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a></p> <p>Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present</p>
<b>lifecycle</b>	<b>object</b>	<p>Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.</p>
<b>livenessProbe</b>	<b>object</b>	<p>Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.</p>
<b>name</b>	<b>string</b>	<p>Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.</p>

Property	Type	Description
<b>ports</b>	<b>array</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

Property	Type	Description
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.12.1.165. .spec.initContainers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.12.1.166. .spec.initContainers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.12.1.167. `.spec.initContainers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources (cpu, memory) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.12.1.168. `.spec.initContainers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a ConfigMap.

#### Type

**object**

#### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.12.1.169. .spec.initContainers[].env[].valueFrom.fieldRef

#### Description

ObjectFieldSelector selects an APIVersioned field of an object.

#### Type

**object**

#### Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.12.1.170. .spec.initContainers[].env[].valueFrom.resourceFieldRef

#### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### Type

**object**

#### Required

- **resource**

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.12.1.171. `.spec.initContainers[].env[].valueFrom.secretKeyRef`

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.12.1.172. `.spec.initContainers[].envFrom`

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.12.1.173. `.spec.initContainers[].envFrom[]`



**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	<p>ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.</p> <p>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.</p>
<b>prefix</b>	<b>string</b>	<p>An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.</p>
<b>secretRef</b>	<b>object</b>	<p>SecretEnvSource selects a Secret to populate the environment variables with.</p> <p>The contents of the target Secret's Data field will represent the key-value pairs as environment variables.</p>

**14.12.1.174. .spec.initContainers[].envFrom[].configMapRef****Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	<p>Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a></p>
<b>optional</b>	<b>boolean</b>	<p>Specify whether the ConfigMap must be defined</p>

### 14.12.1.175. .spec.initContainers[].envFrom[].secretRef

#### Description

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.12.1.176. .spec.initContainers[].lifecycle

#### Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

#### Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.12.1.177. .spec.initContainers[].lifecycle.postStart

#### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

Type  
object

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.12.1.178. .spec.initContainers[].lifecycle.postStart.exec

Description

ExecAction describes a "run in container" action.

Type  
object

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.179. .spec.initContainers[].lifecycle.postStart.httpGet

Description

HTTPGetAction describes an action based on HTTP Get requests.

Type  
object

Required

- port

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.12.1.180. .spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.12.1.181. .spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.182. `.spec.initContainers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.183. `.spec.initContainers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPsocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.12.1.184. .spec.initContainers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.185. .spec.initContainers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.12.1.186. .spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.12.1.187. .spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.188. .spec.initContainers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.189. .spec.initContainers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.



Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.190. .spec.initContainers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.191. .spec.initContainers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.192. .spec.initContainers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.12.1.193. .spec.initContainers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.12.1.194. .spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.195. .spec.initContainers[].livenessProbe.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.196. .spec.initContainers[].ports

##### Description

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <https://github.com/kubernetes/kubernetes/issues/108255>. Cannot be updated.

##### Type

**array**

#### 14.12.1.197. .spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

**object**

**Required**

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

**14.12.1.198. .spec.initContainers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.199. .spec.initContainers[].readinessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.200. .spec.initContainers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.201. .spec.initContainers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.12.1.202. .spec.initContainers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.12.1.203. .spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.204. `.spec.initContainers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.205. `.spec.initContainers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.12.1.206. `.spec.initContainers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.12.1.207. .spec.initContainers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.208. .spec.initContainers[].resources.claims

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.209. .spec.initContainers[].resources.claims[]

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.210. .spec.initContainers[].securityContext

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.12.1.211. .spec.initContainers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities



Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.12.1.212. `.spec.initContainers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.12.1.213. `.spec.initContainers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.12.1.214. .spec.initContainers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.12.1.215. .spec.initContainers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.12.1.216. .spec.initContainers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.12.1.217. .spec.initContainers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.12.1.218. .spec.initContainers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.12.1.219. `.spec.initContainers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.12.1.220. `.spec.initContainers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.12.1.221. .spec.initContainers[].startupProbe.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.12.1.222. .spec.initContainers[].volumeDevices

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.12.1.223. .spec.initContainers[].volumeDevices[]

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type



**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.12.1.224. .spec.initContainers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Cannot be updated.

## Type

**array**

**14.12.1.225. .spec.initContainers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '!'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.12.1.226. .spec.os

##### Description

PodOS defines the OS parameters of a pod.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: <a href="https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration">https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration</a> Clients should expect to handle additional values and treat unrecognized values in this field as os: null

#### 14.12.1.227. .spec.readinessGates

##### Description

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True"  
More info: <https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates>

##### Type

**array**

### 14.12.1.228. .spec.readinessGates[]

#### Description

PodReadinessGate contains the reference to a pod condition

#### Type

**object**

#### Required

- **conditionType**

Property	Type	Description
<b>conditionType</b>	<b>string</b>	ConditionType refers to a condition in the pod's condition list with matching type.

### 14.12.1.229. .spec.resourceClaims

#### Description

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.

This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable.

#### Type

**array**

### 14.12.1.230. .spec.resourceClaims[]

#### Description

PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL.

Property	Type	Description
<b>source</b>	<b>object</b>	ClaimSource describes a reference to a ResourceClaim.  Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.

#### 14.12.1.231. .spec.resourceClaims[].source

##### Description

ClaimSource describes a reference to a ResourceClaim.

Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.

##### Type

**object**

Property	Type	Description
<b>resourceClaimName</b>	<b>string</b>	ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.

Property	Type	Description
<b>resourceClaimTemplateName</b>	<b>string</b>	<p>ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.</p> <p>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The name of the ResourceClaim will be &lt;pod name&gt;-&lt;resource name&gt;, where &lt;resource name&gt; is the PodResourceClaim.Name. Pod validation will reject the pod if the concatenated name is not valid for a ResourceClaim (e.g. too long).</p> <p>An existing ResourceClaim with that name that is not owned by the pod will not be used for the pod to avoid using an unrelated resource by mistake. Scheduling and pod startup are then blocked until the unrelated ResourceClaim is removed.</p> <p>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.</p>

#### 14.12.1.232. .spec.schedulingGates

##### Description

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.

SchedulingGates can only be set at pod creation time, and be removed only afterwards.

This is a beta feature enabled by the PodSchedulingReadiness feature gate.

##### Type

**array**

#### 14.12.1.233. .spec.schedulingGates[]

##### Description

PodSchedulingGate is associated to a Pod to guard its scheduling.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the scheduling gate. Each scheduling gate must have a unique name field.

### 14.12.1.234. .spec.securityContext

#### Description

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

#### Type

**object**

Property	Type	Description
<b>fsGroup</b>	<b>integer</b>	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none"> <li>1. The owning GID will be the FSGroup</li> <li>2. The setgid bit is set (new files created in the volume will be owned by FSGroup)</li> <li>3. The permission bits are OR'd with rw-rw----</li> </ol> <p>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows.</p>

Property	Type	Description
<b>fsGroupChangePolicy</b>	<b>string</b>	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Always"</b> indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior. -  <b>"OnRootMismatch"</b> indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume.</p>
<b>runAsGroup</b>	<b>integer</b>	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.</p>



Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>supplementalGroups</b>	<b>array (integer)</b>	A list of groups applied to the first process run in each container, in addition to the container's primary GID, the fsGroup (if specified), and group memberships defined in the container image for the uid of the container process. If unspecified, no additional groups are added to any container. Note that group memberships defined in the container image for the uid of the container process are still effective, even if they are not included in this list. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>sysctls</b>	<b>array</b>	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.
<b>sysctls[]</b>	<b>object</b>	Sysctl defines a kernel parameter to be set
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.12.1.235. .spec.securityContext.seLinuxOptions

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.12.1.236. .spec.securityContext.seccompProfile

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.12.1.237. .spec.securityContext.sysctls

##### Description

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

##### Type

**array**

#### 14.12.1.238. .spec.securityContext.sysctls[]

##### Description

Sysctl defines a kernel parameter to be set

##### Type

**object**

**Required**

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of a property to set
<b>value</b>	<b>string</b>	Value of a property to set

**14.12.1.239. .spec.securityContext.windowsOptions****Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

**object**

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.

Property	Type	Description
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.12.1.240. .spec.tolerations

##### Description

If specified, the pod's tolerations.

##### Type

**array**

#### 14.12.1.241. .spec.tolerations[]

##### Description

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

##### Type

**object**

Property	Type	Description
<b>effect</b>	<b>string</b>	<p>Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.</p> <p>Possible enum values: -  <b>"NoExecute"</b> Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController. -  <b>"NoSchedule"</b> Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler. -  <b>"PreferNoSchedule"</b> Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler.</p>
<b>key</b>	<b>string</b>	<p>Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.</p>
<b>operator</b>	<b>string</b>	<p>Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.</p> <p>Possible enum values: - <b>"Equal"</b>  - <b>"Exists"</b></p>

Property	Type	Description
<b>tolerationSeconds</b>	<b>integer</b>	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
<b>value</b>	<b>string</b>	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

#### 14.12.1.242. .spec.topologySpreadConstraints

##### Description

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

##### Type

**array**

#### 14.12.1.243. .spec.topologySpreadConstraints[]

##### Description

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

##### Type

**object**

##### Required

- **maxSkew**
- **topologyKey**
- **whenUnsatisfiable**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.
<b>matchLabelKeys</b>	<b>array (string)</b>	<p>MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.</p> <p>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).</p>



Property	Type	Description
<b>maxSkew</b>	<b>integer</b>	<p>MaxSkew describes the degree to which pods may be unevenly distributed. When <b>whenUnsatisfiable=DoNotSchedule</b>, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1.   zone1   zone2   zone3     P P   P P   P   - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <b>whenUnsatisfiable=ScheduleAnyway</b>, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.</p>

Property	Type	Description
<b>minDomains</b>	<b>integer</b>	<p>MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.</p> <p>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2:   zone1   zone2   zone3     P P   P P   P P   The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew.</p> <p>This is a beta field and requires the MinDomainsInPodTopologySpread feature gate to be enabled (enabled by default).</p>

Property	Type	Description
<b>nodeAffinityPolicy</b>	<b>string</b>	<p>NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.</p> <p>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>
<b>nodeTaintsPolicy</b>	<b>string</b>	<p>NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.</p> <p>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>

Property	Type	Description
<b>topologyKey</b>	<b>string</b>	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field.

Property	Type	Description
<b>whenUnsatisfiable</b>	<b>string</b>	<p>WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1:   zone1   zone2   zone3     P P P   P   P   If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <b>more</b> imbalanced. It's a required field.</p> <p>Possible enum values: - <b>"DoNotSchedule"</b> instructs the scheduler not to schedule the pod when constraints are not satisfied. - <b>"ScheduleAnyway"</b> instructs the scheduler to schedule the pod even if constraints are not satisfied.</p>

#### 14.12.1.244. .spec.volumes

##### Description

List of volumes that can be mounted by containers belonging to the pod. More info: <https://kubernetes.io/docs/concepts/storage/volumes>

##### Type

**array**

#### 14.12.1.245. .spec.volumes[]

##### Description

Volume represents a named volume in a pod that may be accessed by any container in the pod.

## Type

**object**

## Required

- **name**

Property	Type	Description
<b>awsElasticBlockStore</b>	<b>object</b>	<p>Represents a Persistent Disk resource in AWS.</p> <p>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.</p>
<b>azureDisk</b>	<b>object</b>	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
<b>azureFile</b>	<b>object</b>	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
<b>cephfs</b>	<b>object</b>	Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.
<b>cinder</b>	<b>object</b>	Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

Property	Type	Description
<b>configMap</b>	<b>object</b>	<p>Adapts a ConfigMap into a volume.</p> <p>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.</p>
<b>csi</b>	<b>object</b>	Represents a source location of a volume to mount, managed by an external CSI driver
<b>downwardAPI</b>	<b>object</b>	DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.
<b>emptyDir</b>	<b>object</b>	Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.
<b>ephemeral</b>	<b>object</b>	Represents an ephemeral volume that is handled by a normal storage driver.
<b>fc</b>	<b>object</b>	Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.
<b>flexVolume</b>	<b>object</b>	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

Property	Type	Description
<b>flocker</b>	<b>object</b>	Represents a Flocker volume mounted by the Flocker agent. One and only one of <code>datasetName</code> and <code>datasetUUID</code> should be set. Flocker volumes do not support ownership management or SELinux relabeling.
<b>gcePersistentDisk</b>	<b>object</b>	Represents a Persistent Disk resource in Google Compute Engine.  A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.
<b>gitRepo</b>	<b>object</b>	Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.  DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an <code>EmptyDir</code> into an <code>InitContainer</code> that clones the repo using <code>git</code> , then mount the <code>EmptyDir</code> into the Pod's container.
<b>glusterfs</b>	<b>object</b>	Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.
<b>hostPath</b>	<b>object</b>	Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.



Property	Type	Description
<b>iscsi</b>	<b>object</b>	Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.
<b>name</b>	<b>string</b>	name of the volume. Must be a DNS_LABEL and unique within the pod. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>nfs</b>	<b>object</b>	Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.
<b>persistentVolumeClaim</b>	<b>object</b>	PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).
<b>photonPersistentDisk</b>	<b>object</b>	Represents a Photon Controller persistent disk resource.
<b>portworxVolume</b>	<b>object</b>	PortworxVolumeSource represents a Portworx volume resource.
<b>projected</b>	<b>object</b>	Represents a projected volume source
<b>quobyte</b>	<b>object</b>	Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

Property	Type	Description
<b>rbd</b>	<b>object</b>	Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.
<b>scaleIO</b>	<b>object</b>	ScaleIOVolumeSource represents a persistent ScaleIO volume
<b>secret</b>	<b>object</b>	Adapts a Secret into a volume.  The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.
<b>storageos</b>	<b>object</b>	Represents a StorageOS persistent volume resource.
<b>vsphereVolume</b>	<b>object</b>	Represents a vSphere volume resource.

#### 14.12.1.246. .spec.volumes[].awsElasticBlockStore

##### Description

Represents a Persistent Disk resource in AWS.

An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumelD**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
<b>readOnly</b>	<b>boolean</b>	readOnly value true will force the readOnly setting in VolumeMounts. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>volumeID</b>	<b>string</b>	volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>

#### 14.12.1.247. .spec.volumes[].azureDisk

##### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **diskName**
- **diskURI**

Property	Type	Description
<b>cachingMode</b>	<b>string</b>	cachingMode is the Host Caching mode: None, Read Only, Read Write.  Possible enum values: - <b>"None"</b> - <b>"ReadOnly"</b> - <b>"ReadWrite"</b>
<b>diskName</b>	<b>string</b>	diskName is the Name of the data disk in the blob storage
<b>diskURI</b>	<b>string</b>	diskURI is the URI of data disk in the blob storage
<b>fsType</b>	<b>string</b>	fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>kind</b>	<b>string</b>	kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared  Possible enum values: - <b>"Dedicated"</b> - <b>"Managed"</b> - <b>"Shared"</b>
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

#### 14.12.1.248. .spec.volumes[].azureFile

##### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **secretName**
- **shareName**

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretName</b>	<b>string</b>	secretName is the name of secret that contains Azure Storage Account Name and Key
<b>shareName</b>	<b>string</b>	shareName is the azure share Name

#### 14.12.1.249. .spec.volumes[].cephfs

##### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**

Property	Type	Description
<b>monitors</b>	<b>array (string)</b>	monitors is Required: Monitors is a collection of Ceph monitors More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>path</b>	<b>string</b>	path is Optional: Used as the mounted root, rather than the full Ceph tree, default is /
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

Property	Type	Description
<b>secretFile</b>	<b>string</b>	secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>user</b>	<b>string</b>	user is optional: User is the rados user name, default is admin More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

#### 14.12.1.250. .spec.volumes[].cephfs.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.12.1.251. .spec.volumes[].cinder

##### Description

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>volumeID</b>	<b>string</b>	volumeID used to identify the volume in cinder. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

#### 14.12.1.252. .spec.volumes[].cinder.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.12.1.253. .spec.volumes[].configMap

##### Description

Adapts a ConfigMap into a volume.

The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

Type  
object

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

#### 14.12.1.254. .spec.volumes[].configMap.items

##### Description



items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

#### Type

**array**

#### 14.12.1.255. `.spec.volumes[].configMap.items[]`

#### Description

Maps a string key to a path within a volume.

#### Type

**object**

#### Required

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.12.1.256. `.spec.volumes[].csi`

#### Description

Represents a source location of a volume to mount, managed by an external CSI driver

#### Type

**object**

## Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
<b>fsType</b>	<b>string</b>	fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
<b>nodePublishSecretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>readOnly</b>	<b>boolean</b>	readOnly specifies a read-only configuration for the volume. Defaults to false (read/write).
<b>volumeAttributes</b>	<b>object (string)</b>	volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**14.12.1.257. .spec.volumes[].csi.nodePublishSecretRef****Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

14.12.1.258. `.spec.volumes[].downwardAPI`**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	Items is a list of downward API volume file
<b>items[]</b>	<b>object</b>	DownwardAPIVolumeFile represents information to create the file containing the pod field

14.12.1.259. `.spec.volumes[].downwardAPI.items`**Description**

Items is a list of downward API volume file

**Type**

**array**

14.12.1.260. `.spec.volumes[].downwardAPI.items[]`**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

**object**

**Required**

**path**

- **path**

Property	Type	Description
<b>fieldRef</b>	<b>object</b>	ObjectFieldSelector selects an APIVersioned field of an object.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
<b>resourceFieldRef</b>	<b>object</b>	ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### 14.12.1.261. .spec.volumes[].downwardAPI.items[].fieldRef

##### Description

ObjectFieldSelector selects an APIVersioned field of an object.

##### Type

**object**

##### Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.12.1.262. `.spec.volumes[].downwardAPI.items[].resourceFieldRef`

##### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

##### Type

**object**

##### Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.12.1.263. `.spec.volumes[].emptyDir`

##### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>medium</b>	<b>string</b>	medium represents what type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

Property	Type	Description
<b>sizeLimit</b>	<b>Quantity</b>	sizeLimit is the total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

#### 14.12.1.264. .spec.volumes[].ephemeral

##### Description

Represents an ephemeral volume that is handled by a normal storage driver.

##### Type

**object**

Property	Type	Description
<b>volumeClaimTemplate</b>	<b>object</b>	PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

#### 14.12.1.265. .spec.volumes[].ephemeral.volumeClaimTemplate

##### Description

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

##### Type

**object**

##### Required

- **spec**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
<b>spec</b>	<b>object</b>	PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

#### 14.12.1.266. .spec.volumes[].ephemeral.volumeClaimTemplate.spec

##### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

##### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the desired access modes the volume should have. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>
<b>dataSource</b>	<b>object</b>	TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

Property	Type	Description
<b>dataSourceRef</b>	<b>object</b>	<p>dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: * While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. * While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. * While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.</p>



Property	Type	Description
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over volumes to consider for binding.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of the StorageClass required by the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1</a>
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>
<b>volumeName</b>	<b>string</b>	volumeName is the binding reference to the PersistentVolume backing this claim.

#### 14.12.1.267. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

##### Description

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

##### Type

**object**

##### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

#### 14.12.1.268. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

##### Description

dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: \* While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. \* While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. \* While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

##### Type

**object**

##### Required

- **kind**
- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced
<b>namespace</b>	<b>string</b>	Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

#### 14.12.1.269. .spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

##### Description

ResourceRequirements describes the compute resource requirements.

##### Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>

Property	Type	Description
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.270. `.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.271. `.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.272. `.spec.volumes[].fc`

##### Description

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	<code>fsType</code> is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>lun</b>	<b>integer</b>	<code>lun</code> is Optional: FC target lun number
<b>readOnly</b>	<b>boolean</b>	<code>readOnly</code> is Optional: Defaults to false (read/write). <code>ReadOnly</code> here will force the <code>ReadOnly</code> setting in <code>VolumeMounts</code> .
<b>targetWWNs</b>	<b>array (string)</b>	<code>targetWWNs</code> is Optional: FC target worldwide names (WWNs)
<b>wwids</b>	<b>array (string)</b>	<code>wwids</code> Optional: FC volume world wide identifiers (wwids) Either <code>wwids</code> or combination of <code>targetWWNs</code> and <code>lun</code> must be set, but not both simultaneously.

#### 14.12.1.273. `.spec.volumes[].flexVolume`

##### Description

`FlexVolume` represents a generic volume resource that is provisioned/attached using an exec based plugin.

##### Type

**object**

## Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume.
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
<b>options</b>	<b>object (string)</b>	options is Optional: this field holds extra command options if any.
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**14.12.1.274. .spec.volumes[].flexVolume.secretRef**

## Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

**14.12.1.275. .spec.volumes[].flocker**

## Description

Represents a Flocker volume mounted by the Flocker agent. One and only one of `datasetName` and `datasetUUID` should be set. Flocker volumes do not support ownership management or SELinux relabeling.

#### Type

**object**

Property	Type	Description
<b>datasetName</b>	<b>string</b>	<code>datasetName</code> is Name of the dataset stored as metadata → name on the dataset for Flocker should be considered as deprecated
<b>datasetUUID</b>	<b>string</b>	<code>datasetUUID</code> is the UUID of the dataset. This is unique identifier of a Flocker dataset

#### 14.12.1.276. `.spec.volumes[].gcePersistentDisk`

##### Description

Represents a Persistent Disk resource in Google Compute Engine.

A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

#### Type

**object**

#### Required

- **pdName**

Property	Type	Description
<b>fsType</b>	<b>string</b>	<code>fsType</code> is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

Property	Type	Description
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>pdName</b>	<b>string</b>	pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

#### 14.12.1.277. .spec.volumes[].gitRepo

##### Description

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.

DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

##### Type

**object**

##### Required

- **repository**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>directory</b>	<b>string</b>	directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
<b>repository</b>	<b>string</b>	repository is the URL
<b>revision</b>	<b>string</b>	revision is the commit hash for the specified revision.

#### 14.12.1.278. .spec.volumes[].glusterfs

##### Description

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **endpoints**
- **path**

Property	Type	Description
<b>endpoints</b>	<b>string</b>	endpoints is the endpoint name that details Glusterfs topology. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>path</b>	<b>string</b>	path is the Glusterfs volume path. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

#### 14.12.1.279. .spec.volumes[].hostPath

##### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>

Property	Type	Description
<b>type</b>	<b>string</b>	<p>type for HostPath Volume Defaults to "" More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a></p> <p>Possible enum values: - "" For backwards compatible, leave it empty if unset - <b>"BlockDevice"</b> A block device must exist at the given path - <b>"CharDevice"</b> A character device must exist at the given path - <b>"Directory"</b> A directory must exist at the given path - <b>"DirectoryOrCreate"</b> If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet. - <b>"File"</b> A file must exist at the given path - <b>"FileOrCreate"</b> If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet. - <b>"Socket"</b> A UNIX socket must exist at the given path</p>

#### 14.12.1.280. .spec.volumes[].iscsi

##### Description

Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **targetPortal**
- **iqn**
- **lun**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>chapAuthDiscovery</b>	<b>boolean</b>	chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication
<b>chapAuthSession</b>	<b>boolean</b>	chapAuthSession defines whether support iSCSI Session CHAP authentication
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#iscsi">https://kubernetes.io/docs/concepts/storage/volumes#iscsi</a>
<b>initiatorName</b>	<b>string</b>	initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <target portal>:<volume name> will be created for the connection.
<b>iqn</b>	<b>string</b>	iqn is the target iSCSI Qualified Name.
<b>iscsiInterface</b>	<b>string</b>	iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
<b>lun</b>	<b>integer</b>	lun represents iSCSI Target Lun number.
<b>portals</b>	<b>array (string)</b>	portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

Property	Type	Description
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>targetPortal</b>	<b>string</b>	targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

#### 14.12.1.281. .spec.volumes[].iscsi.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.12.1.282. .spec.volumes[].nfs

##### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **server**
- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path that is exported by the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>server</b>	<b>string</b>	server is the hostname or IP address of the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

#### 14.12.1.283. .spec.volumes[].persistentVolumeClaim

##### Description

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A

PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

##### Type

**object**

##### Required

- **claimName**

Property	Type	Description
<b>claimName</b>	<b>string</b>	claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims">https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims</a>
<b>readOnly</b>	<b>boolean</b>	readOnly Will force the ReadOnly setting in VolumeMounts. Default false.

#### 14.12.1.284. .spec.volumes[].photonPersistentDisk

##### Description

Represents a Photon Controller persistent disk resource.

##### Type

**object**

**Required**

- **pdID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>pdID</b>	<b>string</b>	pdID is the ID that identifies Photon Controller persistent disk

**14.12.1.285. .spec.volumes[].portworxVolume****Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

**object**

**Required**

- **volumelD**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>volumelD</b>	<b>string</b>	volumelD uniquely identifies a Portworx volume

**14.12.1.286. .spec.volumes[].projected****Description**

Represents a projected volume source

Type  
object

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>sources</b>	<b>array</b>	sources is the list of volume projections
<b>sources[]</b>	<b>object</b>	Projection that may be projected along with other supported volume types

#### 14.12.1.287. .spec.volumes[].projected.sources

##### Description

sources is the list of volume projections

Type  
array

#### 14.12.1.288. .spec.volumes[].projected.sources[]

##### Description

Projection that may be projected along with other supported volume types

Type  
object

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>configMap</b>	<b>object</b>	<p>Adapts a ConfigMap into a projected volume.</p> <p>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.</p>
<b>downwardAPI</b>	<b>object</b>	<p>Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.</p>
<b>secret</b>	<b>object</b>	<p>Adapts a secret into a projected volume.</p> <p>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.</p>
<b>serviceAccountToken</b>	<b>object</b>	<p>ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).</p>

#### 14.12.1.289. .spec.volumes[].projected.sources[].configMap

##### Description

Adapts a ConfigMap into a projected volume.

The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

##### Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

**14.12.1.290. .spec.volumes[].projected.sources[].configMap.items****Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.12.1.291. .spec.volumes[].projected.sources[].configMap.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.12.1.292. .spec.volumes[].projected.sources[].downwardAPI

##### Description

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

##### Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	Items is a list of DownwardAPIVolume file
<b>items[]</b>	<b>object</b>	DownwardAPIVolumeFile represents information to create the file containing the pod field

#### 14.12.1.293. .spec.volumes[].projected.sources[].downwardAPI.items

##### Description

Items is a list of DownwardAPIVolume file

Type

**array**

#### 14.12.1.294. `.spec.volumes[].projected.sources[].downwardAPI.items[]`

Description

DownwardAPIVolumeFile represents information to create the file containing the pod field

Type

**object**

Required

- **path**

Property	Type	Description
<b>fieldRef</b>	<b>object</b>	ObjectFieldSelector selects an APIVersioned field of an object.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the <code>..</code> path. Must be utf-8 encoded. The first item of the relative path must not start with <code>..</code>
<b>resourceFieldRef</b>	<b>object</b>	ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### 14.12.1.295. `.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.12.1.296. `.spec.volumes[].projected.sources[].downwardAPI.items[].resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.12.1.297. `.spec.volumes[].projected.sources[].secret`

Description

Adapts a secret into a projected volume.

The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its key must be defined

#### 14.12.1.298. .spec.volumes[].projected.sources[].secret.items

##### Description

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

##### Type

**array**

#### 14.12.1.299. .spec.volumes[].projected.sources[].secret.items[]

##### Description

Maps a string key to a path within a volume.

##### Type

**object**

##### Required

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.12.1.300. .spec.volumes[].projected.sources[].serviceAccountToken

##### Description

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>audience</b>	<b>string</b>	audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.

Property	Type	Description
<b>expirationSeconds</b>	<b>integer</b>	expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
<b>path</b>	<b>string</b>	path is the path relative to the mount point of the file to project the token into.

#### 14.12.1.301. .spec.volumes[].quobyte

##### Description

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **registry**
- **volume**

Property	Type	Description
<b>group</b>	<b>string</b>	group to map volume access to Default is no group
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.



Property	Type	Description
<b>registry</b>	<b>string</b>	registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
<b>tenant</b>	<b>string</b>	tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
<b>user</b>	<b>string</b>	user to map volume access to Defaults to serviceaccount user
<b>volume</b>	<b>string</b>	volume is a string that references an already created Quobyte volume by name.

#### 14.12.1.302. .spec.volumes[].rbd

##### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**
- **image**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#rbd">https://kubernetes.io/docs/concepts/storage/volumes#rbd</a>

Property	Type	Description
<b>image</b>	<b>string</b>	image is the rados image name. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>keyring</b>	<b>string</b>	keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>monitors</b>	<b>array (string)</b>	monitors is a collection of Ceph monitors. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>pool</b>	<b>string</b>	pool is the rados pool name. Default is rbd. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>user</b>	<b>string</b>	user is the rados user name. Default is admin. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

### 14.12.1.303. .spec.volumes[].rbd.secretRef

#### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.12.1.304. .spec.volumes[].scaleIO

##### Description

ScaleIOVolumeSource represents a persistent ScaleIO volume

##### Type

**object**

##### Required

- **gateway**
- **system**
- **secretRef**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
<b>gateway</b>	<b>string</b>	gateway is the host address of the ScaleIO API Gateway.
<b>protectionDomain</b>	<b>string</b>	protectionDomain is the name of the ScaleIO Protection Domain for the configured storage.
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>sslEnabled</b>	<b>boolean</b>	sslEnabled Flag enable/disable SSL communication with Gateway, default false

Property	Type	Description
<b>storageMode</b>	<b>string</b>	storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
<b>storagePool</b>	<b>string</b>	storagePool is the ScaleIO Storage Pool associated with the protection domain.
<b>system</b>	<b>string</b>	system is the name of the storage system as configured in ScaleIO.
<b>volumeName</b>	<b>string</b>	volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source.

#### 14.12.1.305. .spec.volumes[].scaleIO.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.12.1.306. .spec.volumes[].secret

##### Description

Adapts a Secret into a volume.

The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its keys must be defined
<b>secretName</b>	<b>string</b>	secretName is the name of the secret in the pod's namespace to use. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#secret">https://kubernetes.io/docs/concepts/storage/volumes#secret</a>

### 14.12.1.307. .spec.volumes[].secret.items

#### Description

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys

will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.12.1.308. .spec.volumes[].secret.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

**14.12.1.309. .spec.volumes[].storageos****Description**

Represents a StorageOS persistent volume resource.

**Type****object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>volumeName</b>	<b>string</b>	volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
<b>volumeNamespace</b>	<b>string</b>	volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

#### 14.12.1.310. .spec.volumes[].storageos.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.12.1.311. .spec.volumes[].vsphereVolume

##### Description

Represents a vSphere volume resource.

##### Type

**object**

##### Required

- **volumePath**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>storagePolicyID</b>	<b>string</b>	storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
<b>storagePolicyName</b>	<b>string</b>	storagePolicyName is the storage Policy Based Management (SPBM) profile name.
<b>volumePath</b>	<b>string</b>	volumePath is the path that identifies vSphere volume vmdk

#### 14.12.1.312. .status

##### Description

PodStatus represents information about the status of a pod. Status may trail the actual state of a system, especially if the node that hosts the pod cannot contact the control plane.

##### Type

**object**



Property	Type	Description
<b>conditions</b>	<b>array</b>	Current service state of pod. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions</a>
<b>conditions[]</b>	<b>object</b>	PodCondition contains details for the current condition of this pod.
<b>containerStatuses</b>	<b>array</b>	The list has one entry per container in the manifest. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status</a>
<b>containerStatuses[]</b>	<b>object</b>	ContainerStatus contains details for the current status of this container.
<b>ephemeralContainerStatuses</b>	<b>array</b>	Status for any ephemeral containers that have run in this pod.
<b>ephemeralContainerStatuses []</b>	<b>object</b>	ContainerStatus contains details for the current status of this container.
<b>hostIP</b>	<b>string</b>	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
<b>initContainerStatuses</b>	<b>array</b>	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status</a>
<b>initContainerStatuses[]</b>	<b>object</b>	ContainerStatus contains details for the current status of this container.

Property	Type	Description
<b>message</b>	<b>string</b>	A human readable message indicating details about why the pod is in this condition.
<b>nominatedNodeName</b>	<b>string</b>	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
<b>phase</b>	<b>string</b>	<p>The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values:</p> <p>Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers</p>

Property	Type	Description
		<p>in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod.</p> <p>More info:  <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase</a></p> <p>Possible enum values: - <b>"Failed"</b> means that all containers in the pod have terminated, and at least one container has terminated in a failure (exited with a non-zero exit code or was stopped by the system). - <b>"Pending"</b> means the pod has been accepted by the system, but one or more of the containers has not been started. This includes time before being bound to a node, as well as time spent pulling images onto the host. - <b>"Running"</b> means the pod has been bound to a node and all of the containers have been started. At least one container is still running or is in the process of being restarted. - <b>"Succeeded"</b> means that all containers in the pod have voluntarily terminated with a container exit code of 0, and the system is not going to restart any of these containers. - <b>"Unknown"</b> means that for some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod.</p> <p>Deprecated: It isn't being set since 2015  (74da3b14b0c0f658b3bb8d2def5094686d0e9095)</p>
<b>podIP</b>	<b>string</b>	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.

Property	Type	Description
<b>podIPs</b>	<b>array</b>	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.
<b>podIPs[]</b>	<b>object</b>	IP address information for entries in the (plural) PodIPs field. Each entry includes:  IP: An IP address allocated to the pod. Routable at least within the cluster.
<b>qosClass</b>	<b>string</b>	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-qos/#quality-of-service-classes">https://kubernetes.io/docs/concepts/workloads/pods/pod-qos/#quality-of-service-classes</a>  Possible enum values: - <b>"BestEffort"</b> is the BestEffort qos class. - <b>"Burstable"</b> is the Burstable qos class. - <b>"Guaranteed"</b> is the Guaranteed qos class.
<b>reason</b>	<b>string</b>	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
<b>resize</b>	<b>string</b>	Status of resources resize desired for pod's containers. It is empty if no resources resize is pending. Any changes to container resources will automatically set this to "Proposed"
<b>startTime</b>	<b>Time</b>	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

14.12.1.313. `.status.conditions`

## Description

Current service state of pod. More info: <https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions>

## Type

**array**

14.12.1.314. `.status.conditions[]`

## Description

PodCondition contains details for the current condition of this pod.

## Type

**object**

## Required

- **type**
- **status**

Property	Type	Description
<b>lastProbeTime</b>	<b>Time</b>	Last time we probed the condition.
<b>lastTransitionTime</b>	<b>Time</b>	Last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	Human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	Unique, one-word, CamelCase reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status is the status of the condition. Can be True, False, Unknown. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions</a>
<b>type</b>	<b>string</b>	Type is the type of the condition. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions</a>

### 14.12.1.315. `.status.containerStatuses`

#### Description

The list has one entry per container in the manifest. More info:

<https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status>

#### Type

**array**

### 14.12.1.316. `.status.containerStatuses[]`

#### Description

ContainerStatus contains details for the current status of this container.

#### Type

**object**

#### Required

- **name**
- **ready**
- **restartCount**
- **image**
- **imageID**

Property	Type	Description
<b>allocatedResources</b>	<b>object (Quantity)</b>	AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize.
<b>containerID</b>	<b>string</b>	ContainerID is the ID of the container in the format ' <code>&lt;type&gt;://&lt;container_id&gt;</code> '. Where <code>type</code> is a container runtime identifier, returned from Version call of CRI API (for example "containerd").

Property	Type	Description
<b>image</b>	<b>string</b>	Image is the name of container image that the container is running. The container image may not match the image used in the PodSpec, as it may have been resolved by the runtime. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> .
<b>imageID</b>	<b>string</b>	ImageID is the image ID of the container's image. The image ID may not match the image ID of the image used in the PodSpec, as it may have been resolved by the runtime.
<b>lastState</b>	<b>object</b>	ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.
<b>name</b>	<b>string</b>	Name is a DNS_LABEL representing the unique name of the container. Each container in a pod must have a unique name across all container types. Cannot be updated.
<b>ready</b>	<b>boolean</b>	Ready specifies whether the container is currently passing its readiness check. The value will change as readiness probes keep executing. If no readiness probes are specified, this field defaults to true once the container is fully started (see Started field).  The value is typically used to determine whether a container is ready to accept traffic.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.

Property	Type	Description
<b>restartCount</b>	<b>integer</b>	RestartCount holds the number of times the container has been restarted. Kubelet makes an effort to always increment the value, but there are cases when the state may be lost due to node restarts and then the value may be reset to 0. The value is never negative.
<b>started</b>	<b>boolean</b>	Started indicates whether the container has finished its postStart lifecycle hook and passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. In both cases, startup probes will run again. Is always true when no startupProbe is defined and container is running and has passed the postStart lifecycle hook. The null value must be treated the same as false.
<b>state</b>	<b>object</b>	ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

#### 14.12.1.317. .status.containerStatuses[].lastState

##### Description

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

##### Type

**object**

Property	Type	Description
<b>running</b>	<b>object</b>	ContainerStateRunning is a running state of a container.
<b>terminated</b>	<b>object</b>	ContainerStateTerminated is a terminated state of a container.



Property	Type	Description
<b>waiting</b>	<b>object</b>	ContainerStateWaiting is a waiting state of a container.

#### 14.12.1.318. `.status.containerStatuses[].lastState.running`

##### Description

ContainerStateRunning is a running state of a container.

##### Type

**object**

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which the container was last (re-)started

#### 14.12.1.319. `.status.containerStatuses[].lastState.terminated`

##### Description

ContainerStateTerminated is a terminated state of a container.

##### Type

**object**

##### Required

- **exitCode**

Property	Type	Description
<b>containerID</b>	<b>string</b>	Container's ID in the format '<type>://<container_id>'
<b>exitCode</b>	<b>integer</b>	Exit status from the last termination of the container
<b>finishedAt</b>	<b>Time</b>	Time at which the container last terminated
<b>message</b>	<b>string</b>	Message regarding the last termination of the container
<b>reason</b>	<b>string</b>	(brief) reason from the last termination of the container
<b>signal</b>	<b>integer</b>	Signal from the last termination of the container

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which previous execution of the container started

#### 14.12.1.320. .status.containerStatuses[].lastState.waiting

##### Description

ContainerStateWaiting is a waiting state of a container.

##### Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	Message regarding why the container is not yet running.
<b>reason</b>	<b>string</b>	(brief) reason the container is not yet running.

#### 14.12.1.321. .status.containerStatuses[].resources

##### Description

ResourceRequirements describes the compute resource requirements.

##### Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.

Property	Type	Description
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.322. `.status.containerStatuses[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.323. `.status.containerStatuses[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.324. `.status.containerStatuses[].state`

##### Description

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

##### Type

**object**

Property	Type	Description
<b>running</b>	<b>object</b>	ContainerStateRunning is a running state of a container.
<b>terminated</b>	<b>object</b>	ContainerStateTerminated is a terminated state of a container.
<b>waiting</b>	<b>object</b>	ContainerStateWaiting is a waiting state of a container.

#### 14.12.1.325. `.status.containerStatuses[].state.running`

##### Description

ContainerStateRunning is a running state of a container.

##### Type

**object**

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which the container was last (re-)started

#### 14.12.1.326. `.status.containerStatuses[].state.terminated`

##### Description

ContainerStateTerminated is a terminated state of a container.

Type

object

Required

- `exitCode`

Property	Type	Description
<code>containerID</code>	string	Container's ID in the format '<type>://<container_id>'
<code>exitCode</code>	integer	Exit status from the last termination of the container
<code>finishedAt</code>	Time	Time at which the container last terminated
<code>message</code>	string	Message regarding the last termination of the container
<code>reason</code>	string	(brief) reason from the last termination of the container
<code>signal</code>	integer	Signal from the last termination of the container
<code>startedAt</code>	Time	Time at which previous execution of the container started

#### 14.12.1.327. `.status.containerStatuses[].state.waiting`

Description

ContainerStateWaiting is a waiting state of a container.

Type

object

Property	Type	Description
<code>message</code>	string	Message regarding why the container is not yet running.
<code>reason</code>	string	(brief) reason the container is not yet running.

#### 14.12.1.328. `.status.ephemeralContainerStatuses`

Description

Status for any ephemeral containers that have run in this pod.

Type

**array**

#### 14.12.1.329. .status.ephemeralContainerStatuses[]

Description

ContainerStatus contains details for the current status of this container.

Type

**object**

Required

- **name**
- **ready**
- **restartCount**
- **image**
- **imageID**

Property	Type	Description
<b>allocatedResources</b>	<b>object (Quantity)</b>	AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize.
<b>containerID</b>	<b>string</b>	ContainerID is the ID of the container in the format '<type>://<container_id>'. Where type is a container runtime identifier, returned from Version call of CRI API (for example "containerd").
<b>image</b>	<b>string</b>	Image is the name of container image that the container is running. The container image may not match the image used in the PodSpec, as it may have been resolved by the runtime. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> .

Property	Type	Description
<b>imageID</b>	<b>string</b>	ImageID is the image ID of the container's image. The image ID may not match the image ID of the image used in the PodSpec, as it may have been resolved by the runtime.
<b>lastState</b>	<b>object</b>	ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.
<b>name</b>	<b>string</b>	Name is a DNS_LABEL representing the unique name of the container. Each container in a pod must have a unique name across all container types. Cannot be updated.
<b>ready</b>	<b>boolean</b>	Ready specifies whether the container is currently passing its readiness check. The value will change as readiness probes keep executing. If no readiness probes are specified, this field defaults to true once the container is fully started (see Started field).  The value is typically used to determine whether a container is ready to accept traffic.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>restartCount</b>	<b>integer</b>	RestartCount holds the number of times the container has been restarted. Kubelet makes an effort to always increment the value, but there are cases when the state may be lost due to node restarts and then the value may be reset to 0. The value is never negative.

Property	Type	Description
<b>started</b>	<b>boolean</b>	Started indicates whether the container has finished its postStart lifecycle hook and passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. In both cases, startup probes will run again. Is always true when no startupProbe is defined and container is running and has passed the postStart lifecycle hook. The null value must be treated the same as false.
<b>state</b>	<b>object</b>	ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

#### 14.12.1.330. `.status.ephemeralContainerStatuses[].lastState`

##### Description

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

##### Type

**object**

Property	Type	Description
<b>running</b>	<b>object</b>	ContainerStateRunning is a running state of a container.
<b>terminated</b>	<b>object</b>	ContainerStateTerminated is a terminated state of a container.
<b>waiting</b>	<b>object</b>	ContainerStateWaiting is a waiting state of a container.

#### 14.12.1.331. `.status.ephemeralContainerStatuses[].lastState.running`

##### Description

ContainerStateRunning is a running state of a container.



Type  
object

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which the container was last (re-)started

#### 14.12.1.332. `.status.ephemeralContainerStatuses[].lastState.terminated`

Description

ContainerStateTerminated is a terminated state of a container.

Type

object

Required

- **exitCode**

Property	Type	Description
<b>containerID</b>	<b>string</b>	Container's ID in the format ' <code>&lt;type&gt;://&lt;container_id&gt;</code> '
<b>exitCode</b>	<b>integer</b>	Exit status from the last termination of the container
<b>finishedAt</b>	<b>Time</b>	Time at which the container last terminated
<b>message</b>	<b>string</b>	Message regarding the last termination of the container
<b>reason</b>	<b>string</b>	(brief) reason from the last termination of the container
<b>signal</b>	<b>integer</b>	Signal from the last termination of the container
<b>startedAt</b>	<b>Time</b>	Time at which previous execution of the container started

#### 14.12.1.333. `.status.ephemeralContainerStatuses[].lastState.waiting`

Description

ContainerStateWaiting is a waiting state of a container.

Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	Message regarding why the container is not yet running.
<b>reason</b>	<b>string</b>	(brief) reason the container is not yet running.

**14.12.1.334. .status.ephemeralContainerStatuses[].resources****Description**

ResourceRequirements describes the compute resource requirements.

**Type****object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	<p>Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a></p>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.335. `.status.ephemeralContainerStatuses[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.336. `.status.ephemeralContainerStatuses[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.337. `.status.ephemeralContainerStatuses[].state`

##### Description

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

### Type

**object**

Property	Type	Description
<b>running</b>	<b>object</b>	ContainerStateRunning is a running state of a container.
<b>terminated</b>	<b>object</b>	ContainerStateTerminated is a terminated state of a container.
<b>waiting</b>	<b>object</b>	ContainerStateWaiting is a waiting state of a container.

#### 14.12.1.338. .status.ephemeralContainerStatuses[].state.running

### Description

ContainerStateRunning is a running state of a container.

### Type

**object**

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which the container was last (re-)started

#### 14.12.1.339. .status.ephemeralContainerStatuses[].state.terminated

### Description

ContainerStateTerminated is a terminated state of a container.

### Type

**object**

### Required

- **exitCode**

Property	Type	Description
<b>containerID</b>	<b>string</b>	Container's ID in the format '<type>://<container_id>'
<b>exitCode</b>	<b>integer</b>	Exit status from the last termination of the container

Property	Type	Description
<b>finishedAt</b>	<b>Time</b>	Time at which the container last terminated
<b>message</b>	<b>string</b>	Message regarding the last termination of the container
<b>reason</b>	<b>string</b>	(brief) reason from the last termination of the container
<b>signal</b>	<b>integer</b>	Signal from the last termination of the container
<b>startedAt</b>	<b>Time</b>	Time at which previous execution of the container started

#### 14.12.1.340. `.status.ephemeralContainerStatuses[].state.waiting`

##### Description

ContainerStateWaiting is a waiting state of a container.

##### Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	Message regarding why the container is not yet running.
<b>reason</b>	<b>string</b>	(brief) reason the container is not yet running.

#### 14.12.1.341. `.status.initContainerStatuses`

##### Description

The list has one entry per init container in the manifest. The most recent successful init container will have `ready = true`, the most recently started container will have `startTime` set. More info:

<https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status>

##### Type

**array**

#### 14.12.1.342. `.status.initContainerStatuses[]`

##### Description

ContainerStatus contains details for the current status of this container.

##### Type

**object**

## Required

- **name**
- **ready**
- **restartCount**
- **image**
- **imageID**

Property	Type	Description
<b>allocatedResources</b>	<b>object (Quantity)</b>	AllocatedResources represents the compute resources allocated for this container by the node. Kubelet sets this value to Container.Resources.Requests upon successful pod admission and after successfully admitting desired pod resize.
<b>containerID</b>	<b>string</b>	ContainerID is the ID of the container in the format '<type>://<container_id>'. Where type is a container runtime identifier, returned from Version call of CRI API (for example "containerd").
<b>image</b>	<b>string</b>	Image is the name of container image that the container is running. The container image may not match the image used in the PodSpec, as it may have been resolved by the runtime. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> .
<b>imageID</b>	<b>string</b>	ImageID is the image ID of the container's image. The image ID may not match the image ID of the image used in the PodSpec, as it may have been resolved by the runtime.

Property	Type	Description
<b>lastState</b>	<b>object</b>	ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.
<b>name</b>	<b>string</b>	Name is a DNS_LABEL representing the unique name of the container. Each container in a pod must have a unique name across all container types. Cannot be updated.
<b>ready</b>	<b>boolean</b>	Ready specifies whether the container is currently passing its readiness check. The value will change as readiness probes keep executing. If no readiness probes are specified, this field defaults to true once the container is fully started (see Started field).  The value is typically used to determine whether a container is ready to accept traffic.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>restartCount</b>	<b>integer</b>	RestartCount holds the number of times the container has been restarted. Kubelet makes an effort to always increment the value, but there are cases when the state may be lost due to node restarts and then the value may be reset to 0. The value is never negative.

Property	Type	Description
<b>started</b>	<b>boolean</b>	Started indicates whether the container has finished its postStart lifecycle hook and passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. In both cases, startup probes will run again. Is always true when no startupProbe is defined and container is running and has passed the postStart lifecycle hook. The null value must be treated the same as false.
<b>state</b>	<b>object</b>	ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

#### 14.12.1.343. `.status.initContainerStatuses[].lastState`

##### Description

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

##### Type

**object**

Property	Type	Description
<b>running</b>	<b>object</b>	ContainerStateRunning is a running state of a container.
<b>terminated</b>	<b>object</b>	ContainerStateTerminated is a terminated state of a container.
<b>waiting</b>	<b>object</b>	ContainerStateWaiting is a waiting state of a container.

#### 14.12.1.344. `.status.initContainerStatuses[].lastState.running`

##### Description

ContainerStateRunning is a running state of a container.



Type  
object

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which the container was last (re-)started

#### 14.12.1.345. `.status.initContainerStatuses[].lastState.terminated`

Description

ContainerStateTerminated is a terminated state of a container.

Type

object

Required

- **exitCode**

Property	Type	Description
<b>containerID</b>	<b>string</b>	Container's ID in the format ' <code>&lt;type&gt;://&lt;container_id&gt;</code> '
<b>exitCode</b>	<b>integer</b>	Exit status from the last termination of the container
<b>finishedAt</b>	<b>Time</b>	Time at which the container last terminated
<b>message</b>	<b>string</b>	Message regarding the last termination of the container
<b>reason</b>	<b>string</b>	(brief) reason from the last termination of the container
<b>signal</b>	<b>integer</b>	Signal from the last termination of the container
<b>startedAt</b>	<b>Time</b>	Time at which previous execution of the container started

#### 14.12.1.346. `.status.initContainerStatuses[].lastState.waiting`

Description

ContainerStateWaiting is a waiting state of a container.

Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	Message regarding why the container is not yet running.
<b>reason</b>	<b>string</b>	(brief) reason the container is not yet running.

**14.12.1.347. .status.initContainerStatuses[].resources****Description**

ResourceRequirements describes the compute resource requirements.

**Type****object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	<p>Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a></p>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.12.1.348. `.status.initContainerStatuses[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.12.1.349. `.status.initContainerStatuses[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.12.1.350. `.status.initContainerStatuses[].state`

##### Description

ContainerState holds a possible state of container. Only one of its members may be specified. If none of them is specified, the default one is ContainerStateWaiting.

### Type

**object**

Property	Type	Description
<b>running</b>	<b>object</b>	ContainerStateRunning is a running state of a container.
<b>terminated</b>	<b>object</b>	ContainerStateTerminated is a terminated state of a container.
<b>waiting</b>	<b>object</b>	ContainerStateWaiting is a waiting state of a container.

#### 14.12.1.351. .status.initContainerStatuses[].state.running

### Description

ContainerStateRunning is a running state of a container.

### Type

**object**

Property	Type	Description
<b>startedAt</b>	<b>Time</b>	Time at which the container was last (re-)started

#### 14.12.1.352. .status.initContainerStatuses[].state.terminated

### Description

ContainerStateTerminated is a terminated state of a container.

### Type

**object**

### Required

- **exitCode**

Property	Type	Description
<b>containerID</b>	<b>string</b>	Container's ID in the format '<type>://<container_id>'
<b>exitCode</b>	<b>integer</b>	Exit status from the last termination of the container

Property	Type	Description
<b>finishedAt</b>	<b>Time</b>	Time at which the container last terminated
<b>message</b>	<b>string</b>	Message regarding the last termination of the container
<b>reason</b>	<b>string</b>	(brief) reason from the last termination of the container
<b>signal</b>	<b>integer</b>	Signal from the last termination of the container
<b>startedAt</b>	<b>Time</b>	Time at which previous execution of the container started

#### 14.12.1.353. `.status.initContainerStatuses[].state.waiting`

##### Description

ContainerStateWaiting is a waiting state of a container.

##### Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	Message regarding why the container is not yet running.
<b>reason</b>	<b>string</b>	(brief) reason the container is not yet running.

#### 14.12.1.354. `.status.podIPs`

##### Description

podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

##### Type

**array**

#### 14.12.1.355. `.status.podIPs[]`

##### Description

IP address information for entries in the (plural) PodIPs field. Each entry includes:

- IP: An IP address allocated to the pod. Routable at least within the cluster.

Type  
object

Property	Type	Description
<b>ip</b>	<b>string</b>	ip is an IP address (IPv4 or IPv6) assigned to the pod

### 14.12.2. API endpoints

The following API endpoints are available:

- **/api/v1/pods**
  - **GET**: list or watch objects of kind Pod
- **/api/v1/watch/pods**
  - **GET**: watch individual changes to a list of Pod. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/pods**
  - **DELETE**: delete collection of Pod
  - **GET**: list or watch objects of kind Pod
  - **POST**: create a Pod
- **/api/v1/watch/namespaces/{namespace}/pods**
  - **GET**: watch individual changes to a list of Pod. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/pods/{name}**
  - **DELETE**: delete a Pod
  - **GET**: read the specified Pod
  - **PATCH**: partially update the specified Pod
  - **PUT**: replace the specified Pod
- **/api/v1/namespaces/{namespace}/pods/{name}/log**
  - **GET**: read log of the specified Pod
- **/api/v1/watch/namespaces/{namespace}/pods/{name}**
  - **GET**: watch changes to an object of kind Pod. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/namespaces/{namespace}/pods/{name}/status**
  - **GET**: read status of the specified Pod

- **PATCH**: partially update status of the specified Pod
- **PUT**: replace status of the specified Pod
- **/api/v1/namespaces/{namespace}/pods/{name}/ephemeralcontainers**
  - **GET**: read ephemeralcontainers of the specified Pod
  - **PATCH**: partially update ephemeralcontainers of the specified Pod
  - **PUT**: replace ephemeralcontainers of the specified Pod

#### 14.12.2.1. /api/v1/pods

Table 14.289. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
GET

**Description**

list or watch objects of kind Pod

**Table 14.290. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">PodList</a> schema
401 - Unauthorized	Empty

**14.12.2.2. /api/v1/watch/pods****Table 14.291. Global query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

## HTTP method

**GET**

## Description

watch individual changes to a list of Pod. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.292. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 14.12.2.3. /api/v1/namespaces/{namespace}/pods

Table 14.293. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.294. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**DELETE**

## Description

delete collection of Pod

Table 14.295. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>



Table 14.296. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.297. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Pod

Table 14.298. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.299. HTTP responses

HTTP code	Response body
200 - OK	<b>PodList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a Pod

**Table 14.300. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.301. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Pod</b> schema	

**Table 14.302. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Pod</b> schema

HTTP code	Response body
201 - Created	<b>Pod</b> schema
202 - Accepted	<b>Pod</b> schema
401 - Unauthorized	Empty

#### 14.12.2.4. /api/v1/watch/namespaces/{namespace}/pods

Table 14.303. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.304. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Pod. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.305. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.12.2.5. /api/v1/namespaces/{namespace}/pods/{name}

Table 14.306. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Pod
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.307. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a Pod

Table 14.308. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.309. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.310. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema
202 - Accepted	<b>Pod</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Pod

Table 14.311. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Pod

**Table 14.312. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.313. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.314. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema
201 - Created	<b>Pod</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Pod

Table 14.315. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.316. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Pod</b> schema	

Table 14.317. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema
201 - Created	<b>Pod</b> schema
401 - Unauthorized	Empty

#### 14.12.2.6. /api/v1/namespaces/{namespace}/pods/{name}/log

Table 14.318. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Pod
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.319. Global query parameters

Parameter	Type	Description
<b>container</b>	<b>string</b>	The container for which to stream logs. Defaults to only container if there is one container in the pod.
<b>follow</b>	<b>boolean</b>	Follow the log stream of the pod. Defaults to false.
<b>insecureSkipTLSVerifyBackend</b>	<b>boolean</b>	<code>insecureSkipTLSVerifyBackend</code> indicates that the apiserver should not confirm the validity of the serving certificate of the backend it is connecting to. This will make the HTTPS connection between the apiserver and the backend insecure. This means the apiserver cannot verify the log data it is receiving came from the real kubelet. If the kubelet is configured to verify the apiserver's TLS credentials, it does not mean the connection to the real kubelet is vulnerable to a man in the middle attack (e.g. an attacker could not intercept the actual log data coming from the real kubelet).
<b>limitBytes</b>	<b>integer</b>	If set, the number of bytes to read from the server before terminating the log output. This may not display a complete final line of logging, and may return slightly more or slightly less than the specified limit.
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>previous</b>	<b>boolean</b>	Return previous terminated container logs. Defaults to false.
<b>sinceSeconds</b>	<b>integer</b>	A relative time in seconds before the current time from which to show logs. If this value precedes the time a pod was started, only logs since the pod start will be returned. If this value is in the future, no logs will be returned. Only one of <code>sinceSeconds</code> or <code>sinceTime</code> may be specified.
<b>tailLines</b>	<b>integer</b>	If set, the number of lines from the end of the logs to show. If not specified, logs are shown from the creation of the container or <code>sinceSeconds</code> or <code>sinceTime</code>
<b>timestamps</b>	<b>boolean</b>	If true, add an RFC3339 or RFC3339Nano timestamp at the beginning of every line of log output. Defaults to false.

HTTP method

GET

**Description**

read log of the specified Pod

**Table 14.320. HTTP responses**

HTTP code	Response body
200 - OK	<b>string</b>
401 - Unauthorized	Empty

**14.12.2.7. /api/v1/watch/namespaces/{namespace}/pods/{name}****Table 14.321. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Pod
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

**Table 14.322. Global query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Pod. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.323. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.12.2.8. /api/v1/namespaces/{namespace}/pods/{name}/status

Table 14.324. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Pod
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.325. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Pod

Table 14.326. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">Pod</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Pod

Table 14.327. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.328. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.329. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema
201 - Created	<b>Pod</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Pod

**Table 14.330. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.331. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Pod</b> schema	

Table 14.332. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Pod</b> schema
201 - Created	<b>Pod</b> schema
401 - Unauthorized	Empty

#### 14.12.2.9. /api/v1/namespaces/{namespace}/pods/{name}/ephemeralcontainers

Table 14.333. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Pod
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.334. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read ephemeralcontainers of the specified Pod

Table 14.335. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Pod</b> schema
401 - Unauthorized	Empty

#### HTTP method

**PATCH****Description**

partially update ephemeralcontainers of the specified Pod

**Table 14.336. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

**Table 14.337. Body parameters**



Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.338. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema
201 - Created	<b>Pod</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace ephemeralcontainers of the specified Pod

Table 14.339. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.340. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Pod</b> schema	

Table 14.341. HTTP responses

HTTP code	Response body
200 - OK	<b>Pod</b> schema
201 - Created	<b>Pod</b> schema
401 - Unauthorized	Empty

## 14.13. PODTEMPLATE [V1]

### Description

PodTemplate describes a template for creating copies of a predefined pod.

### Type

**object**

### 14.13.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>template</b>	<b>object</b>	PodTemplateSpec describes the data a pod should have when created from a template

#### 14.13.1.1. .template

##### Description

PodTemplateSpec describes the data a pod should have when created from a template

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PodSpec is a description of a pod.

#### 14.13.1.2. .template.spec

##### Description

PodSpec is a description of a pod.

##### Type

**object**

##### Required

- **containers**

Property	Type	Description
<b>activeDeadlineSeconds</b>	<b>integer</b>	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
<b>affinity</b>	<b>object</b>	Affinity is a group of affinity scheduling rules.
<b>automountServiceAccountToken</b>	<b>boolean</b>	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
<b>containers</b>	<b>array</b>	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
<b>containers[]</b>	<b>object</b>	A single application container that you want to run within a pod.

Property	Type	Description
<b>dnsConfig</b>	<b>object</b>	PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.
<b>dnsPolicy</b>	<b>string</b>	<p>Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.</p> <p>Possible enum values: - <b>"ClusterFirst"</b> indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"ClusterFirstWithHostNet"</b> indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"Default"</b> indicates that the pod should use the default (as determined by kubelet) DNS settings. - <b>"None"</b> indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig.</p>
<b>enableServiceLinks</b>	<b>boolean</b>	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.

Property	Type	Description
<b>ephemeralContainers</b>	<b>array</b>	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's <code>ephemeralcontainers</code> subresource.
<b>ephemeralContainers[]</b>	<b>object</b>	<p>An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation.</p> <p>To add an ephemeral container, use the <code>ephemeralcontainers</code> subresource of an existing Pod. Ephemeral containers may not be removed or restarted.</p>
<b>hostAliases</b>	<b>array</b>	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
<b>hostAliases[]</b>	<b>object</b>	HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.
<b>hostIPC</b>	<b>boolean</b>	Use the host's ipc namespace. Optional: Default to false.

Property	Type	Description
<b>hostNetwork</b>	<b>boolean</b>	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
<b>hostPID</b>	<b>boolean</b>	Use the host's pid namespace. Optional: Default to false.
<b>hostUsers</b>	<b>boolean</b>	Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new users is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature.
<b>hostname</b>	<b>string</b>	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
<b>imagePullSecrets</b>	<b>array</b>	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod">https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod</a>
<b>imagePullSecrets[]</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

Property	Type	Description
<b>initContainers</b>	<b>array</b>	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/init-containers/">https://kubernetes.io/docs/concepts/workloads/pods/init-containers/</a>
<b>initContainers[]</b>	<b>object</b>	A single application container that you want to run within a pod.
<b>nodeName</b>	<b>string</b>	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
<b>nodeSelector</b>	<b>object (string)</b>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a>



Property	Type	Description
<b>os</b>	<b>object</b>	PodOS defines the OS parameters of a pod.
<b>overhead</b>	<b>object (Quantity)</b>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md">https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md</a>
<b>preemptionPolicy</b>	<b>string</b>	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.  Possible enum values: - <b>"Never"</b> means that pod never preempts other pods with lower priority. - <b>"PreemptLowerPriority"</b> means that pod can preempt other pods with lower priority.
<b>priority</b>	<b>integer</b>	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

Property	Type	Description
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
<b>readinessGates</b>	<b>array</b>	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: <a href="https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates">https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates</a>
<b>readinessGates[]</b>	<b>object</b>	PodReadinessGate contains the reference to a pod condition
<b>resourceClaims</b>	<b>array</b>	ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable.
<b>resourceClaims[]</b>	<b>object</b>	PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

Property	Type	Description
<b>restartPolicy</b>	<b>string</b>	<p>Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy</a></p> <p>Possible enum values: - <b>"Always"</b> - <b>"Never"</b> - <b>"OnFailure"</b></p>
<b>runtimeClassName</b>	<b>string</b>	<p>RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class">https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class</a></p>
<b>schedulerName</b>	<b>string</b>	<p>If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.</p>
<b>schedulingGates</b>	<b>array</b>	<p>SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.</p> <p>SchedulingGates can only be set at pod creation time, and be removed only afterwards.</p> <p>This is a beta feature enabled by the PodSchedulingReadiness feature gate.</p>

Property	Type	Description
<b>schedulingGates[]</b>	<b>object</b>	PodSchedulingGate is associated to a Pod to guard its scheduling.
<b>securityContext</b>	<b>object</b>	PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.
<b>serviceAccount</b>	<b>string</b>	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
<b>serviceAccountName</b>	<b>string</b>	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: <a href="https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/">https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/</a>

Property	Type	Description
<b>setHostnameAsFQDN</b>	<b>boolean</b>	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
<b>shareProcessNamespace</b>	<b>boolean</b>	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
<b>subdomain</b>	<b>string</b>	If specified, the fully qualified Pod hostname will be "<hostname>.<subdomain>.<pod namespace>.svc.<cluster domain>". If not specified, the pod will not have a domainname at all.

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
<b>tolerations</b>	<b>array</b>	If specified, the pod's tolerations.
<b>tolerations[]</b>	<b>object</b>	The pod this Tolerant is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.
<b>topologySpreadConstraints</b>	<b>array</b>	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
<b>topologySpreadConstraints[]</b>	<b>object</b>	TopologySpreadConstraint specifies how to spread matching pods among the given topology.
<b>volumes</b>	<b>array</b>	List of volumes that can be mounted by containers belonging to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes">https://kubernetes.io/docs/concepts/storage/volumes</a>
<b>volumes[]</b>	<b>object</b>	Volume represents a named volume in a pod that may be accessed by any container in the pod.

### 14.13.1.3. .template.spec.affinity

#### Description

Affinity is a group of affinity scheduling rules.

#### Type

**object**

Property	Type	Description
<b>nodeAffinity</b>	<b>object</b>	Node affinity is a group of node affinity scheduling rules.
<b>podAffinity</b>	<b>object</b>	Pod affinity is a group of inter pod affinity scheduling rules.
<b>podAntiAffinity</b>	<b>object</b>	Pod anti affinity is a group of inter pod anti affinity scheduling rules.

### 14.13.1.4. .template.spec.affinity.nodeAffinity

#### Description

Node affinity is a group of node affinity scheduling rules.

#### Type

**object**

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>object</b>	A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### 14.13.1.5. `.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution`

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, `requiredDuringScheduling` affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding `matchExpressions`; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.13.1.6. `.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecution`

##### Description

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

##### Type

**object**

##### Required

- **weight**
- **preference**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>preference</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.
<b>weight</b>	<b>integer</b>	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

#### 14.13.1.7. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecu

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.13.1.8. .template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecu

##### Description

A list of node selector requirements by node's labels.

##### Type

**array**

### 14.13.1.9. `.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecu`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.13.1.10. `.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExec`

#### Description

A list of node selector requirements by node's fields.

#### Type

**array**

### 14.13.1.11. `.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuringExecu`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 14.13.1.12. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu

##### Description

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### Type

**object**

#### Required

- **nodeSelectorTerms**

Property	Type	Description
<b>nodeSelectorTerms</b>	<b>array</b>	Required. A list of node selector terms. The terms are ORed.
<b>nodeSelectorTerms[]</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

#### 14.13.1.13. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu

##### Description

Required. A list of node selector terms. The terms are ORed.

##### Type

**array**

#### 14.13.1.14. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.13.1.15. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu

##### Description

A list of node selector requirements by node's labels.

Type

**array**

#### 14.13.1.16. `.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu`

Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

Type

**object**

Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 14.13.1.17. `.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu`

Description

A list of node selector requirements by node's fields.

Type

**array**

### 14.13.1.18. .template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecu

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.13.1.19. .template.spec.affinity.podAffinity

#### Description

Pod affinity is a group of inter pod affinity scheduling rules.

#### Type

**object**

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

#### 14.13.1.20. .template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.13.1.21. .template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecution

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

##### Type

**object**

##### Required

- **weight**
- **podAffinityTerm**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>podAffinityTerm</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

#### 14.13.1.22. .template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringExecu

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.

Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.13.1.23. `.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecut`

##### Description

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

##### Type

**array**

#### 14.13.1.24. `.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringExecut`

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.13.1.25. .template.spec.affinity.podAntiAffinity

##### Description

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

#### 14.13.1.26. .template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringE

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.13.1.27. .template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringE

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

##### Type

**object**

##### Required

- **weight**
- **podAffinityTerm**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>podAffinityTerm</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

#### 14.13.1.28. .template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDuringE

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} matches all namespaces.

Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.13.1.29. `.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringEx`

##### Description

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

##### Type

**array**

#### 14.13.1.30. `.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDuringEx`

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.13.1.31. .template.spec.containers

##### Description

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

##### Type

**array**

#### 14.13.1.32. .template.spec.containers[]

##### Description

A single application container that you want to run within a pod.



Type

**object**

Required

- **name**

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references <code>\$(VAR_NAME)</code> are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>\$</code> are reduced to a single <code>\$</code>, which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. "<code>\$(VAR_NAME)</code>" will produce the string literal "<code>\$(VAR_NAME)</code>". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>
<b>env</b>	<b>array</b>	List of environment variables to set in the container. Cannot be updated.
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps

Property	Type	Description
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
<b>imagePullPolicy</b>	<b>string</b>	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a>  Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present
<b>lifecycle</b>	<b>object</b>	Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.
<b>livenessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
<b>ports</b>	<b>array</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

Property	Type	Description
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.13.1.33. .template.spec.containers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.13.1.34. .template.spec.containers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.13.1.35. `.template.spec.containers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources ( <code>cpu</code> , <code>memory</code> ) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.13.1.36. `.template.spec.containers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a `ConfigMap`.



Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.13.1.37. `.template.spec.containers[].env[].valueFrom.fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.13.1.38. `.template.spec.containers[].env[].valueFrom.resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.13.1.39. .template.spec.containers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.13.1.40. .template.spec.containers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.13.1.41. .template.spec.containers[].envFrom[]

##### Description

EnvFromSource represents the source of a set of ConfigMaps

#### Type

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>object</b>	SecretEnvSource selects a Secret to populate the environment variables with.  The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### 14.13.1.42. .template.spec.containers[].envFrom[].configMapRef

##### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  
The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 14.13.1.43. .template.spec.containers[].envFrom[].secretRef

## Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.13.1.44. .template.spec.containers[].lifecycle

## Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

## Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.13.1.45. .template.spec.containers[].lifecycle.postStart

## Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

**14.13.1.46. .template.spec.containers[].lifecycle.postStart.exec****Description**

ExecAction describes a "run in container" action.

**Type**

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**14.13.1.47. .template.spec.containers[].lifecycle.postStart.httpGet****Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

**object**

**Required**

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.13.1.48. .template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.13.1.49. .template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.50. `.template.spec.containers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.51. `.template.spec.containers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.13.1.52. .template.spec.containers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.53. .template.spec.containers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.



Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.13.1.54. .template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.13.1.55. .template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.56. .template.spec.containers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.57. .template.spec.containers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.58. .template.spec.containers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.59. .template.spec.containers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.60. .template.spec.containers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.13.1.61. .template.spec.containers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.13.1.62. .template.spec.containers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.63. `.template.spec.containers[].livenessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.64. `.template.spec.containers[].ports`

##### Description

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <https://github.com/kubernetes/kubernetes/issues/108255>. Cannot be updated.

##### Type

**array**

#### 14.13.1.65. `.template.spec.containers[].ports[]`

**Description**

ContainerPort represents a network port in a single container.

**Type**

**object**

**Required**

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

**14.13.1.66. .template.spec.containers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

**object**



Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.67. .template.spec.containers[].readinessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.68. .template.spec.containers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.69. .template.spec.containers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.13.1.70. .template.spec.containers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.13.1.71. .template.spec.containers[].readinessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.72. `.template.spec.containers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.73. `.template.spec.containers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.13.1.74. `.template.spec.containers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.13.1.75. .template.spec.containers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.13.1.76. .template.spec.containers[].resources.claims

##### Description

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.13.1.77. .template.spec.containers[].resources.claims[]

##### Description

ResourceClaim references one entry in PodSpec.ResourceClaims.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container.

#### 14.13.1.78. .template.spec.containers[].securityContext

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.



Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.13.1.79. .template.spec.containers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities

Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.13.1.80. .template.spec.containers[].securityContext.seLinuxOptions

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.13.1.81. .template.spec.containers[].securityContext.seccompProfile

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.13.1.82. .template.spec.containers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

### 14.13.1.83. .template.spec.containers[].startupProbe

#### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

#### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.84. .template.spec.containers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.85. .template.spec.containers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.86. .template.spec.containers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**



## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.13.1.87. `.template.spec.containers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.13.1.88. `.template.spec.containers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.89. `.template.spec.containers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.90. `.template.spec.containers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.13.1.91. `.template.spec.containers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.13.1.92. .template.spec.containers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Cannot be updated.

## Type

**array**

**14.13.1.93. .template.spec.containers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '!'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: - <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). - <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.13.1.94. .template.spec.dnsConfig

##### Description

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

##### Type

**object**

Property	Type	Description
<b>nameservers</b>	<b>array (string)</b>	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
<b>options</b>	<b>array</b>	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
<b>options[]</b>	<b>object</b>	PodDNSConfigOption defines DNS resolver options of a pod.
<b>searches</b>	<b>array (string)</b>	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

#### 14.13.1.95. .template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

**array**

**14.13.1.96. .template.spec.dnsConfig.options[]****Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Required.
<b>value</b>	<b>string</b>	

**14.13.1.97. .template.spec.ephemeralContainers****Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

**array**

**14.13.1.98. .template.spec.ephemeralContainers[]****Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

**object**

**Required**

- **name**

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Endpoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array</b>	<p>List of environment variables to set in the container. Cannot be updated.</p>

Property	Type	Description
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a>
<b>imagePullPolicy</b>	<b>string</b>	<p>Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a></p> <p>Possible enum values: -  <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present</p>



Property	Type	Description
<b>lifecycle</b>	<b>object</b>	Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.
<b>livenessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>name</b>	<b>string</b>	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
<b>ports</b>	<b>array</b>	Ports are not allowed for ephemeral containers.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.

Property	Type	Description
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

Property	Type	Description
<b>targetContainerName</b>	<b>string</b>	<p>If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.</p> <p>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined.</p>
<b>terminationMessagePath</b>	<b>string</b>	<p>Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to <code>/dev/termination-log</code>. Cannot be updated.</p>

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.13.1.99. .template.spec.ephemeralContainers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.13.1.100. .template.spec.ephemeralContainers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.13.1.101. `.template.spec.ephemeralContainers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources ( <code>cpu</code> , <code>memory</code> ) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.13.1.102. `.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a `ConfigMap`.

Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.13.1.103. `.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.13.1.104. `.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.13.1.105. .template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.13.1.106. .template.spec.ephemeralContainers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.13.1.107. .template.spec.ephemeralContainers[].envFrom[]

##### Description



EnvFromSource represents the source of a set of ConfigMaps

### Type

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>object</b>	SecretEnvSource selects a Secret to populate the environment variables with.  The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### 14.13.1.108. .template.spec.ephemeralContainers[].envFrom[].configMapRef

### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  
The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 14.13.1.109. .template.spec.ephemeralContainers[].envFrom[].secretRef

## Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.13.1.110. .template.spec.ephemeralContainers[].lifecycle

## Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

## Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.13.1.111. .template.spec.ephemeralContainers[].lifecycle.postStart

## Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

**14.13.1.112. .template.spec.ephemeralContainers[].lifecycle.postStart.exec****Description**

ExecAction describes a "run in container" action.

**Type**

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**14.13.1.113. .template.spec.ephemeralContainers[].lifecycle.postStart.httpGet****Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

**object**

**Required**

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.13.1.114. .template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.13.1.115. .template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.116. `.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.117. `.template.spec.ephemeralContainers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.13.1.118. .template.spec.ephemeralContainers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.119. .template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.13.1.120. .template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.13.1.121. .template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.122. .template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.123. .template.spec.ephemeralContainers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.



Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.124. `.template.spec.ephemeralContainers[].livenessProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.125. .template.spec.ephemeralContainers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.126. .template.spec.ephemeralContainers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.13.1.127. `.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.13.1.128. `.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.129. `.template.spec.ephemeralContainers[].livenessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.130. `.template.spec.ephemeralContainers[].ports`

##### Description

Ports are not allowed for ephemeral containers.

##### Type

**array**

#### 14.13.1.131. `.template.spec.ephemeralContainers[].ports[]`

##### Description

ContainerPort represents a network port in a single container.

##### Type

**object**

## Required

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.

**14.13.1.132. .template.spec.ephemeralContainers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type****object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.133. `.template.spec.ephemeralContainers[].readinessProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.134. .template.spec.ephemeralContainers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.135. .template.spec.ephemeralContainers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.13.1.136. .template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.13.1.137. .template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.138. `.template.spec.ephemeralContainers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.139. `.template.spec.ephemeralContainers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.13.1.140. `.template.spec.ephemeralContainers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.13.1.141. .template.spec.ephemeralContainers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.13.1.142. `.template.spec.ephemeralContainers[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.13.1.143. `.template.spec.ephemeralContainers[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.13.1.144. `.template.spec.ephemeralContainers[].securityContext`

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.13.1.145. .template.spec.ephemeralContainers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities



Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.13.1.146. .template.spec.ephemeralContainers[].securityContext.seLinuxOptions

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.13.1.147. .template.spec.ephemeralContainers[].securityContext.seccompProfile

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.13.1.148. .template.spec.ephemeralContainers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.13.1.149. .template.spec.ephemeralContainers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.150. .template.spec.ephemeralContainers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.151. .template.spec.ephemeralContainers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.152. .template.spec.ephemeralContainers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.13.1.153. `.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.13.1.154. `.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.155. `.template.spec.ephemeralContainers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.156. `.template.spec.ephemeralContainers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.13.1.157. `.template.spec.ephemeralContainers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type



**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.13.1.158. .template.spec.ephemeralContainers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

## Type

**array**

**14.13.1.159. .template.spec.ephemeralContainers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '.'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). -  <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.13.1.160. .template.spec.hostAliases

##### Description

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.

##### Type

**array**

#### 14.13.1.161. .template.spec.hostAliases[]

##### Description

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

##### Type

**object**

Property	Type	Description
<b>hostnames</b>	<b>array (string)</b>	Hostnames for the above IP address.
<b>ip</b>	<b>string</b>	IP address of the host file entry.

#### 14.13.1.162. .template.spec.imagePullSecrets

##### Description

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info:

<https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod>

##### Type

**array**

#### 14.13.1.163. .template.spec.imagePullSecrets[]

## Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

### 14.13.1.164. .template.spec.initContainers

## Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <https://kubernetes.io/docs/concepts/workloads/pods/init-containers/>

## Type

**array**

### 14.13.1.165. .template.spec.initContainers[]

## Description

A single application container that you want to run within a pod.

## Type

**object**

## Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The container image's CMD is used if this is not provided. Variable references <code>\$(VAR_NAME)</code> are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single <code>\$</code>, which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. "<code>\$(VAR_NAME)</code>" will produce the string literal "<code>\$(VAR_NAME)</code>". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Endpoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references <code>\$(VAR_NAME)</code> are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single <code>\$</code>, which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. "<code>\$(VAR_NAME)</code>" will produce the string literal "<code>\$(VAR_NAME)</code>". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array</b>	<p>List of environment variables to set in the container. Cannot be updated.</p>

Property	Type	Description
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.

Property	Type	Description
<b>imagePullPolicy</b>	<b>string</b>	<p>Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a></p> <p>Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present</p>
<b>lifecycle</b>	<b>object</b>	<p>Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.</p>
<b>livenessProbe</b>	<b>object</b>	<p>Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.</p>
<b>name</b>	<b>string</b>	<p>Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.</p>

Property	Type	Description
<b>ports</b>	<b>array</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.



Property	Type	Description
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.13.1.166. .template.spec.initContainers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.13.1.167. .template.spec.initContainers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.13.1.168. `.template.spec.initContainers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources (cpu, memory) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.13.1.169. `.template.spec.initContainers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a ConfigMap.

Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.13.1.170. .template.spec.initContainers[].env[].valueFrom.fieldRef

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.13.1.171. .template.spec.initContainers[].env[].valueFrom.resourceFieldRef

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.13.1.172. .template.spec.initContainers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.13.1.173. .template.spec.initContainers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.13.1.174. .template.spec.initContainers[].envFrom[]

**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	<p>ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.</p> <p>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.</p>
<b>prefix</b>	<b>string</b>	<p>An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.</p>
<b>secretRef</b>	<b>object</b>	<p>SecretEnvSource selects a Secret to populate the environment variables with.</p> <p>The contents of the target Secret's Data field will represent the key-value pairs as environment variables.</p>

**14.13.1.175. .template.spec.initContainers[].envFrom[].configMapRef****Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	<p>Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a></p>
<b>optional</b>	<b>boolean</b>	<p>Specify whether the ConfigMap must be defined</p>

### 14.13.1.176. .template.spec.initContainers[].envFrom[].secretRef

#### Description

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.13.1.177. .template.spec.initContainers[].lifecycle

#### Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

#### Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.13.1.178. .template.spec.initContainers[].lifecycle.postStart

#### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.



Type  
object

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.13.1.179. .template.spec.initContainers[].lifecycle.postStart.exec

Description

ExecAction describes a "run in container" action.

Type  
object

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.180. .template.spec.initContainers[].lifecycle.postStart.httpGet

Description

HTTPGetAction describes an action based on HTTP Get requests.

Type  
object

Required

- port

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.13.1.181. .template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.13.1.182. .template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.183. `.template.spec.initContainers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.184. `.template.spec.initContainers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCP socket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.13.1.185. .template.spec.initContainers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.186. .template.spec.initContainers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.13.1.187. .template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.13.1.188. .template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.189. .template.spec.initContainers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.190. .template.spec.initContainers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.191. .template.spec.initContainers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.192. .template.spec.initContainers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.193. .template.spec.initContainers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.13.1.194. .template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.13.1.195. .template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.196. .template.spec.initContainers[].livenessProbe.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.197. .template.spec.initContainers[].ports

##### Description

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <https://github.com/kubernetes/kubernetes/issues/108255>. Cannot be updated.

##### Type

**array**

#### 14.13.1.198. .template.spec.initContainers[].ports[]

**Description**

ContainerPort represents a network port in a single container.

**Type**

**object**

**Required**

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

**14.13.1.199. .template.spec.initContainers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.200. `.template.spec.initContainers[].readinessProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.201. .template.spec.initContainers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.202. .template.spec.initContainers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.13.1.203. .template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.13.1.204. .template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required



- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.205. `.template.spec.initContainers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.206. `.template.spec.initContainers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.13.1.207. `.template.spec.initContainers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.13.1.208. .template.spec.initContainers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.13.1.209. .template.spec.initContainers[].resources.claims

##### Description

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.13.1.210. .template.spec.initContainers[].resources.claims[]

##### Description

ResourceClaim references one entry in PodSpec.ResourceClaims.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container.

#### 14.13.1.211. .template.spec.initContainers[].securityContext

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

## Type

**object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.13.1.212. .template.spec.initContainers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities

Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.13.1.213. `.template.spec.initContainers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.13.1.214. `.template.spec.initContainers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.13.1.215. .template.spec.initContainers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.13.1.216. .template.spec.initContainers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.13.1.217. `.template.spec.initContainers[].startupProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.13.1.218. .template.spec.initContainers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.13.1.219. .template.spec.initContainers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.13.1.220. .template.spec.initContainers[].startupProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.13.1.221. .template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.13.1.222. `.template.spec.initContainers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.13.1.223. `.template.spec.initContainers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.13.1.224. `.template.spec.initContainers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.13.1.225. .template.spec.initContainers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Cannot be updated.

## Type

**array**

**14.13.1.226. .template.spec.initContainers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '!'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).



Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

### 14.13.1.227. .template.spec.os

#### Description

PodOS defines the OS parameters of a pod.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: <a href="https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration">https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration</a> Clients should expect to handle additional values and treat unrecognized values in this field as os: null

### 14.13.1.228. .template.spec.readinessGates

#### Description

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True"  
More info: <https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates>

#### Type

**array**

### 14.13.1.229. .template.spec.readinessGates[]

#### Description

PodReadinessGate contains the reference to a pod condition

#### Type

**object**

#### Required

- **conditionType**

Property	Type	Description
<b>conditionType</b>	<b>string</b>	ConditionType refers to a condition in the pod's condition list with matching type.

### 14.13.1.230. .template.spec.resourceClaims

#### Description

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.

This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable.

#### Type

**array**

### 14.13.1.231. .template.spec.resourceClaims[]

#### Description

PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL.

Property	Type	Description
<b>source</b>	<b>object</b>	<p>ClaimSource describes a reference to a ResourceClaim.</p> <p>Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.</p>

#### 14.13.1.232. .template.spec.resourceClaims[].source

##### Description

ClaimSource describes a reference to a ResourceClaim.

Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.

##### Type

**object**

Property	Type	Description
<b>resourceClaimName</b>	<b>string</b>	ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.

Property	Type	Description
<b>resourceClaimTemplateName</b>	<b>string</b>	<p>ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.</p> <p>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The name of the ResourceClaim will be &lt;pod name&gt;-&lt;resource name&gt;, where &lt;resource name&gt; is the PodResourceClaim.Name. Pod validation will reject the pod if the concatenated name is not valid for a ResourceClaim (e.g. too long).</p> <p>An existing ResourceClaim with that name that is not owned by the pod will not be used for the pod to avoid using an unrelated resource by mistake. Scheduling and pod startup are then blocked until the unrelated ResourceClaim is removed.</p> <p>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.</p>

### 14.13.1.233. .template.spec.schedulingGates

#### Description

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.

SchedulingGates can only be set at pod creation time, and be removed only afterwards.

This is a beta feature enabled by the PodSchedulingReadiness feature gate.

#### Type

**array**

### 14.13.1.234. .template.spec.schedulingGates[]

#### Description

PodSchedulingGate is associated to a Pod to guard its scheduling.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the scheduling gate. Each scheduling gate must have a unique name field.

### 14.13.1.235. .template.spec.securityContext

#### Description

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

#### Type

**object**

Property	Type	Description
<b>fsGroup</b>	<b>integer</b>	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none"> <li>1. The owning GID will be the FSGroup</li> <li>2. The setgid bit is set (new files created in the volume will be owned by FSGroup)</li> <li>3. The permission bits are OR'd with rw-rw----</li> </ol> <p>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows.</p>

Property	Type	Description
<b>fsGroupChangePolicy</b>	<b>string</b>	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Always"</b> indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior. -  <b>"OnRootMismatch"</b> indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume.</p>
<b>runAsGroup</b>	<b>integer</b>	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.</p>

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>supplementalGroups</b>	<b>array (integer)</b>	A list of groups applied to the first process run in each container, in addition to the container's primary GID, the fsGroup (if specified), and group memberships defined in the container image for the uid of the container process. If unspecified, no additional groups are added to any container. Note that group memberships defined in the container image for the uid of the container process are still effective, even if they are not included in this list. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>sysctls</b>	<b>array</b>	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.
<b>sysctls[]</b>	<b>object</b>	Sysctl defines a kernel parameter to be set
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.13.1.236. .template.spec.securityContext.seLinuxOptions

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.13.1.237. .template.spec.securityContext.seccompProfile

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**



Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.13.1.238. .template.spec.securityContext.sysctls

##### Description

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

##### Type

**array**

#### 14.13.1.239. .template.spec.securityContext.sysctls[]

##### Description

Sysctl defines a kernel parameter to be set

##### Type

**object**

**Required**

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of a property to set
<b>value</b>	<b>string</b>	Value of a property to set

**14.13.1.240. .template.spec.securityContext.windowsOptions****Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

**object**

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.

Property	Type	Description
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.13.1.241. .template.spec.tolerations

##### Description

If specified, the pod's tolerations.

##### Type

**array**

#### 14.13.1.242. .template.spec.tolerations[]

##### Description

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

##### Type

**object**

Property	Type	Description
<b>effect</b>	<b>string</b>	<p>Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.</p> <p>Possible enum values: -  <b>"NoExecute"</b> Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController. -  <b>"NoSchedule"</b> Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler. -  <b>"PreferNoSchedule"</b> Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler.</p>
<b>key</b>	<b>string</b>	<p>Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.</p>
<b>operator</b>	<b>string</b>	<p>Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.</p> <p>Possible enum values: - <b>"Equal"</b>  - <b>"Exists"</b></p>

Property	Type	Description
<b>tolerationSeconds</b>	<b>integer</b>	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
<b>value</b>	<b>string</b>	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

#### 14.13.1.243. `.template.spec.topologySpreadConstraints`

##### Description

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

##### Type

**array**

#### 14.13.1.244. `.template.spec.topologySpreadConstraints[]`

##### Description

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

##### Type

**object**

##### Required

- **maxSkew**
- **topologyKey**
- **whenUnsatisfiable**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.
<b>matchLabelKeys</b>	<b>array (string)</b>	<p>MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.</p> <p>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).</p>

Property	Type	Description
<b>maxSkew</b>	<b>integer</b>	<p>MaxSkew describes the degree to which pods may be unevenly distributed. When <b>whenUnsatisfiable=DoNotSchedule</b>, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1.   zone1   zone2   zone3     P P   P P   P   - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <b>whenUnsatisfiable=ScheduleAnyway</b>, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.</p>

Property	Type	Description
<b>minDomains</b>	<b>integer</b>	<p>MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.</p> <p>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2:   zone1   zone2   zone3     P P   P P   P P   The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew.</p> <p>This is a beta field and requires the MinDomainsInPodTopologySpread feature gate to be enabled (enabled by default).</p>



Property	Type	Description
<b>nodeAffinityPolicy</b>	<b>string</b>	<p>NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.</p> <p>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>
<b>nodeTaintsPolicy</b>	<b>string</b>	<p>NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.</p> <p>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>

Property	Type	Description
<b>topologyKey</b>	<b>string</b>	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field.

Property	Type	Description
<b>whenUnsatisfiable</b>	<b>string</b>	<p>WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1:   zone1   zone2   zone3     P P P   P   P   If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <b>more</b> imbalanced. It's a required field.</p> <p>Possible enum values: - <b>"DoNotSchedule"</b> instructs the scheduler not to schedule the pod when constraints are not satisfied. - <b>"ScheduleAnyway"</b> instructs the scheduler to schedule the pod even if constraints are not satisfied.</p>

#### 14.13.1.245. .template.spec.volumes

##### Description

List of volumes that can be mounted by containers belonging to the pod. More info: <https://kubernetes.io/docs/concepts/storage/volumes>

##### Type

**array**

#### 14.13.1.246. .template.spec.volumes[]

##### Description

Volume represents a named volume in a pod that may be accessed by any container in the pod.

## Type

**object**

## Required

- **name**

Property	Type	Description
<b>awsElasticBlockStore</b>	<b>object</b>	<p>Represents a Persistent Disk resource in AWS.</p> <p>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.</p>
<b>azureDisk</b>	<b>object</b>	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
<b>azureFile</b>	<b>object</b>	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
<b>cephfs</b>	<b>object</b>	Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.
<b>cinder</b>	<b>object</b>	Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

Property	Type	Description
<b>configMap</b>	<b>object</b>	<p>Adapts a ConfigMap into a volume.</p> <p>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.</p>
<b>csi</b>	<b>object</b>	Represents a source location of a volume to mount, managed by an external CSI driver
<b>downwardAPI</b>	<b>object</b>	DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.
<b>emptyDir</b>	<b>object</b>	Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.
<b>ephemeral</b>	<b>object</b>	Represents an ephemeral volume that is handled by a normal storage driver.
<b>fc</b>	<b>object</b>	Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.
<b>flexVolume</b>	<b>object</b>	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

Property	Type	Description
<b>flocker</b>	<b>object</b>	Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.
<b>gcePersistentDisk</b>	<b>object</b>	Represents a Persistent Disk resource in Google Compute Engine.  A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.
<b>gitRepo</b>	<b>object</b>	Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.  DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
<b>glusterfs</b>	<b>object</b>	Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.
<b>hostPath</b>	<b>object</b>	Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

Property	Type	Description
<b>iscsi</b>	<b>object</b>	Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.
<b>name</b>	<b>string</b>	name of the volume. Must be a DNS_LABEL and unique within the pod. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>nfs</b>	<b>object</b>	Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.
<b>persistentVolumeClaim</b>	<b>object</b>	PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).
<b>photonPersistentDisk</b>	<b>object</b>	Represents a Photon Controller persistent disk resource.
<b>portworxVolume</b>	<b>object</b>	PortworxVolumeSource represents a Portworx volume resource.
<b>projected</b>	<b>object</b>	Represents a projected volume source
<b>quobyte</b>	<b>object</b>	Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

Property	Type	Description
<b>rbd</b>	<b>object</b>	Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.
<b>scaleIO</b>	<b>object</b>	ScaleIOVolumeSource represents a persistent ScaleIO volume
<b>secret</b>	<b>object</b>	Adapts a Secret into a volume.  The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.
<b>storageos</b>	<b>object</b>	Represents a StorageOS persistent volume resource.
<b>vsphereVolume</b>	<b>object</b>	Represents a vSphere volume resource.

#### 14.13.1.247. .template.spec.volumes[].awsElasticBlockStore

##### Description

Represents a Persistent Disk resource in AWS.

An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumelD**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
<b>readOnly</b>	<b>boolean</b>	readOnly value true will force the readOnly setting in VolumeMounts. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>volumeID</b>	<b>string</b>	volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>

#### 14.13.1.248. .template.spec.volumes[].azureDisk

##### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **diskName**
- **diskURI**

Property	Type	Description
<b>cachingMode</b>	<b>string</b>	cachingMode is the Host Caching mode: None, Read Only, Read Write.  Possible enum values: - <b>"None"</b> - <b>"ReadOnly"</b> - <b>"ReadWrite"</b>
<b>diskName</b>	<b>string</b>	diskName is the Name of the data disk in the blob storage
<b>diskURI</b>	<b>string</b>	diskURI is the URI of data disk in the blob storage
<b>fsType</b>	<b>string</b>	fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>kind</b>	<b>string</b>	kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared  Possible enum values: - <b>"Dedicated"</b> - <b>"Managed"</b> - <b>"Shared"</b>
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

#### 14.13.1.249. .template.spec.volumes[].azureFile

##### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **secretName**
- **shareName**

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretName</b>	<b>string</b>	secretName is the name of secret that contains Azure Storage Account Name and Key
<b>shareName</b>	<b>string</b>	shareName is the azure share Name

#### 14.13.1.250. .template.spec.volumes[].cephfs

##### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**

Property	Type	Description
<b>monitors</b>	<b>array (string)</b>	monitors is Required: Monitors is a collection of Ceph monitors More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>path</b>	<b>string</b>	path is Optional: Used as the mounted root, rather than the full Ceph tree, default is /
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

Property	Type	Description
<b>secretFile</b>	<b>string</b>	secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>user</b>	<b>string</b>	user is optional: User is the rados user name, default is admin More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

#### 14.13.1.251. .template.spec.volumes[].cephfs.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.13.1.252. .template.spec.volumes[].cinder

##### Description

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>volumeID</b>	<b>string</b>	volumeID used to identify the volume in cinder. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

#### 14.13.1.253. .template.spec.volumes[].cinder.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.13.1.254. .template.spec.volumes[].configMap

##### Description

Adapts a ConfigMap into a volume.

The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

Type  
object

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

14.13.1.255. `.template.spec.volumes[].configMap.items`

Description

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

#### Type

**array**

### 14.13.1.256. .template.spec.volumes[].configMap.items[]

#### Description

Maps a string key to a path within a volume.

#### Type

**object**

#### Required

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

### 14.13.1.257. .template.spec.volumes[].csi

#### Description

Represents a source location of a volume to mount, managed by an external CSI driver

#### Type

**object**

## Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
<b>fsType</b>	<b>string</b>	fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
<b>nodePublishSecretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>readOnly</b>	<b>boolean</b>	readOnly specifies a read-only configuration for the volume. Defaults to false (read/write).
<b>volumeAttributes</b>	<b>object (string)</b>	volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**14.13.1.258. .template.spec.volumes[].csi.nodePublishSecretRef**

## Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>



### 14.13.1.259. .template.spec.volumes[].downwardAPI

#### Description

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

#### Type

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	Items is a list of downward API volume file
<b>items[]</b>	<b>object</b>	DownwardAPIVolumeFile represents information to create the file containing the pod field

### 14.13.1.260. .template.spec.volumes[].downwardAPI.items

#### Description

Items is a list of downward API volume file

#### Type

**array**

### 14.13.1.261. .template.spec.volumes[].downwardAPI.items[]

#### Description

DownwardAPIVolumeFile represents information to create the file containing the pod field

#### Type

**object**

#### Required

**path**

- **path**

Property	Type	Description
<b>fieldRef</b>	<b>object</b>	ObjectFieldSelector selects an APIVersioned field of an object.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
<b>resourceFieldRef</b>	<b>object</b>	ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### 14.13.1.262. .template.spec.volumes[].downwardAPI.items[].fieldRef

##### Description

ObjectFieldSelector selects an APIVersioned field of an object.

##### Type

**object**

##### Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.13.1.263. .template.spec.volumes[].downwardAPI.items[].resourceFieldRef

##### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

##### Type

**object**

##### Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.13.1.264. .template.spec.volumes[].emptyDir

##### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>medium</b>	<b>string</b>	medium represents what type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

Property	Type	Description
<b>sizeLimit</b>	<b>Quantity</b>	sizeLimit is the total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

#### 14.13.1.265. .template.spec.volumes[].ephemeral

##### Description

Represents an ephemeral volume that is handled by a normal storage driver.

##### Type

**object**

Property	Type	Description
<b>volumeClaimTemplate</b>	<b>object</b>	PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

#### 14.13.1.266. .template.spec.volumes[].ephemeral.volumeClaimTemplate

##### Description

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

##### Type

**object**

##### Required

- **spec**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
<b>spec</b>	<b>object</b>	PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

#### 14.13.1.267. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

##### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

##### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the desired access modes the volume should have. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>
<b>dataSource</b>	<b>object</b>	TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

Property	Type	Description
<b>dataSourceRef</b>	<b>object</b>	<p>dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: * While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. * While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. * While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.</p>

Property	Type	Description
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over volumes to consider for binding.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of the StorageClass required by the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1</a>
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>
<b>volumeName</b>	<b>string</b>	volumeName is the binding reference to the PersistentVolume backing this claim.

#### 14.13.1.268. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

##### Description

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

##### Type

**object**

##### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

### 14.13.1.269. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef

#### Description

`dataSourceRef` specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a `PersistentVolumeClaim` object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the `dataSource` field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in `dataSourceRef`, both fields (`dataSource` and `dataSourceRef`) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in `dataSourceRef`, `dataSource` isn't set to the same value and must be empty. There are three important differences between `dataSource` and `dataSourceRef`: \* While `dataSource` only allows two specific types of objects, `dataSourceRef` allows any non-core object, as well as `PersistentVolumeClaim` objects. \* While `dataSource` ignores disallowed values (dropping them), `dataSourceRef` preserves all values, and generates an error if a disallowed value is specified. \* While `dataSource` only allows local objects, `dataSourceRef` allows objects in any namespaces. (Beta) Using this field requires the `AnyVolumeDataSource` feature gate to be enabled. (Alpha) Using the namespace field of `dataSourceRef` requires the `CrossNamespaceVolumeDataSource` feature gate to be enabled.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced
<b>namespace</b>	<b>string</b>	Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

#### 14.13.1.270. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

##### Description

ResourceRequirements describes the compute resource requirements.

##### Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>

Property	Type	Description
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.13.1.271. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims

##### Description

Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container. This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.13.1.272. .template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.claims

##### Description

ResourceClaim references one entry in PodSpec.ResourceClaims.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container.

#### 14.13.1.273. .template.spec.volumes[].fc

##### Description

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>lun</b>	<b>integer</b>	lun is Optional: FC target lun number
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>targetWWNs</b>	<b>array (string)</b>	targetWWNs is Optional: FC target worldwide names (WWNs)
<b>wwids</b>	<b>array (string)</b>	wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

#### 14.13.1.274. .template.spec.volumes[].flexVolume

##### Description

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

##### Type

**object**

## Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume.
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
<b>options</b>	<b>object (string)</b>	options is Optional: this field holds extra command options if any.
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**14.13.1.275. .template.spec.volumes[].flexVolume.secretRef****Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

**14.13.1.276. .template.spec.volumes[].flocker****Description**

Represents a Flocker volume mounted by the Flocker agent. One and only one of `datasetName` and `datasetUUID` should be set. Flocker volumes do not support ownership management or SELinux relabeling.

#### Type

**object**

Property	Type	Description
<b>datasetName</b>	<b>string</b>	<code>datasetName</code> is Name of the dataset stored as metadata → name on the dataset for Flocker should be considered as deprecated
<b>datasetUUID</b>	<b>string</b>	<code>datasetUUID</code> is the UUID of the dataset. This is unique identifier of a Flocker dataset

#### 14.13.1.277. `.template.spec.volumes[].gcePersistentDisk`

##### Description

Represents a Persistent Disk resource in Google Compute Engine.

A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

#### Type

**object**

#### Required

- `pdName`

Property	Type	Description
<b>fsType</b>	<b>string</b>	<code>fsType</code> is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

Property	Type	Description
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>pdName</b>	<b>string</b>	pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

#### 14.13.1.278. .template.spec.volumes[].gitRepo

##### Description

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.

DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

##### Type

**object**

##### Required

- **repository**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>directory</b>	<b>string</b>	directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
<b>repository</b>	<b>string</b>	repository is the URL
<b>revision</b>	<b>string</b>	revision is the commit hash for the specified revision.

#### 14.13.1.279. .template.spec.volumes[].glusterfs

##### Description

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **endpoints**
- **path**

Property	Type	Description
<b>endpoints</b>	<b>string</b>	endpoints is the endpoint name that details Glusterfs topology. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>path</b>	<b>string</b>	path is the Glusterfs volume path. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

#### 14.13.1.280. .template.spec.volumes[].hostPath

##### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>



Property	Type	Description
<b>type</b>	<b>string</b>	<p>type for HostPath Volume Defaults to "" More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a></p> <p>Possible enum values: - "" For backwards compatible, leave it empty if unset - <b>"BlockDevice"</b> A block device must exist at the given path - <b>"CharDevice"</b> A character device must exist at the given path - <b>"Directory"</b> A directory must exist at the given path - <b>"DirectoryOrCreate"</b> If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet. - <b>"File"</b> A file must exist at the given path - <b>"FileOrCreate"</b> If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet. - <b>"Socket"</b> A UNIX socket must exist at the given path</p>

#### 14.13.1.281. .template.spec.volumes[].iscsi

##### Description

Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **targetPortal**
- **iqn**
- **lun**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>chapAuthDiscovery</b>	<b>boolean</b>	chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication
<b>chapAuthSession</b>	<b>boolean</b>	chapAuthSession defines whether support iSCSI Session CHAP authentication
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#iscsi">https://kubernetes.io/docs/concepts/storage/volumes#iscsi</a>
<b>initiatorName</b>	<b>string</b>	initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <target portal>:<volume name> will be created for the connection.
<b>iqn</b>	<b>string</b>	iqn is the target iSCSI Qualified Name.
<b>iscsiInterface</b>	<b>string</b>	iscsiInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
<b>lun</b>	<b>integer</b>	lun represents iSCSI Target Lun number.
<b>portals</b>	<b>array (string)</b>	portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

Property	Type	Description
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>targetPortal</b>	<b>string</b>	targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

#### 14.13.1.282. .template.spec.volumes[].iscsi.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.13.1.283. .template.spec.volumes[].nfs

##### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **server**
- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path that is exported by the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>server</b>	<b>string</b>	server is the hostname or IP address of the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

#### 14.13.1.284. .template.spec.volumes[].persistentVolumeClaim

##### Description

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A

PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

##### Type

**object**

##### Required

- **claimName**

Property	Type	Description
<b>claimName</b>	<b>string</b>	claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims">https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims</a>
<b>readOnly</b>	<b>boolean</b>	readOnly Will force the ReadOnly setting in VolumeMounts. Default false.

#### 14.13.1.285. .template.spec.volumes[].photonPersistentDisk

##### Description

Represents a Photon Controller persistent disk resource.

##### Type

**object**

**Required**

- **pdID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>pdID</b>	<b>string</b>	pdID is the ID that identifies Photon Controller persistent disk

**14.13.1.286. .template.spec.volumes[].portworxVolume****Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

**object**

**Required**

- **volumelD**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>volumelD</b>	<b>string</b>	volumelD uniquely identifies a Portworx volume

**14.13.1.287. .template.spec.volumes[].projected****Description**

Represents a projected volume source

## Type

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>sources</b>	<b>array</b>	sources is the list of volume projections
<b>sources[]</b>	<b>object</b>	Projection that may be projected along with other supported volume types

14.13.1.288. `.template.spec.volumes[].projected.sources`

## Description

sources is the list of volume projections

## Type

**array**14.13.1.289. `.template.spec.volumes[].projected.sources[]`

## Description

Projection that may be projected along with other supported volume types

## Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>configMap</b>	<b>object</b>	<p>Adapts a ConfigMap into a projected volume.</p> <p>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.</p>
<b>downwardAPI</b>	<b>object</b>	<p>Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.</p>
<b>secret</b>	<b>object</b>	<p>Adapts a secret into a projected volume.</p> <p>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.</p>
<b>serviceAccountToken</b>	<b>object</b>	<p>ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).</p>

#### 14.13.1.290. `.template.spec.volumes[].projected.sources[].configMap`

##### Description

Adapts a ConfigMap into a projected volume.

The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

##### Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

**14.13.1.291. .template.spec.volumes[].projected.sources[].configMap.items****Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.13.1.292. .template.spec.volumes[].projected.sources[].configMap.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**



- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.13.1.293. .template.spec.volumes[].projected.sources[].downwardAPI

##### Description

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

##### Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	Items is a list of DownwardAPIVolume file
<b>items[]</b>	<b>object</b>	DownwardAPIVolumeFile represents information to create the file containing the pod field

#### 14.13.1.294. .template.spec.volumes[].projected.sources[].downwardAPI.items

##### Description

Items is a list of DownwardAPIVolume file

Type

**array**

#### 14.13.1.295. .template.spec.volumes[].projected.sources[].downwardAPI.items[]

Description

DownwardAPIVolumeFile represents information to create the file containing the pod field

Type

**object**

Required

- **path**

Property	Type	Description
<b>fieldRef</b>	<b>object</b>	ObjectFieldSelector selects an APIVersioned field of an object.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
<b>resourceFieldRef</b>	<b>object</b>	ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### 14.13.1.296. .template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.13.1.297. .template.spec.volumes[].projected.sources[].downwardAPI.items[].resourceField

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.13.1.298. .template.spec.volumes[].projected.sources[].secret

Description

Adapts a secret into a projected volume.

The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its key must be defined

#### 14.13.1.299. .template.spec.volumes[].projected.sources[].secret.items

##### Description

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

##### Type

**array**

#### 14.13.1.300. .template.spec.volumes[].projected.sources[].secret.items[]

##### Description

Maps a string key to a path within a volume.

##### Type

**object**

##### Required

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.13.1.301. .template.spec.volumes[].projected.sources[].serviceAccountToken

##### Description

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>audience</b>	<b>string</b>	audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.

Property	Type	Description
<b>expirationSeconds</b>	<b>integer</b>	expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
<b>path</b>	<b>string</b>	path is the path relative to the mount point of the file to project the token into.

#### 14.13.1.302. .template.spec.volumes[].quobyte

##### Description

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **registry**
- **volume**

Property	Type	Description
<b>group</b>	<b>string</b>	group to map volume access to Default is no group
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.

Property	Type	Description
<b>registry</b>	<b>string</b>	registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
<b>tenant</b>	<b>string</b>	tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
<b>user</b>	<b>string</b>	user to map volume access to Defaults to serviceaccount user
<b>volume</b>	<b>string</b>	volume is a string that references an already created Quobyte volume by name.

#### 14.13.1.303. .template.spec.volumes[].rbd

##### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**
- **image**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#rbd">https://kubernetes.io/docs/concepts/storage/volumes#rbd</a>

Property	Type	Description
<b>image</b>	<b>string</b>	image is the rados image name. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>keyring</b>	<b>string</b>	keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>monitors</b>	<b>array (string)</b>	monitors is a collection of Ceph monitors. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>pool</b>	<b>string</b>	pool is the rados pool name. Default is rbd. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>user</b>	<b>string</b>	user is the rados user name. Default is admin. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

#### 14.13.1.304. .template.spec.volumes[].rbd.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**



Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.13.1.305. .template.spec.volumes[].scaleIO

##### Description

ScaleIOVolumeSource represents a persistent ScaleIO volume

##### Type

**object**

##### Required

- **gateway**
- **system**
- **secretRef**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
<b>gateway</b>	<b>string</b>	gateway is the host address of the ScaleIO API Gateway.
<b>protectionDomain</b>	<b>string</b>	protectionDomain is the name of the ScaleIO Protection Domain for the configured storage.
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>sslEnabled</b>	<b>boolean</b>	sslEnabled Flag enable/disable SSL communication with Gateway, default false

Property	Type	Description
<b>storageMode</b>	<b>string</b>	storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
<b>storagePool</b>	<b>string</b>	storagePool is the ScaleIO Storage Pool associated with the protection domain.
<b>system</b>	<b>string</b>	system is the name of the storage system as configured in ScaleIO.
<b>volumeName</b>	<b>string</b>	volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source.

#### 14.13.1.306. .template.spec.volumes[].scaleIO.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.13.1.307. .template.spec.volumes[].secret

##### Description

Adapts a Secret into a volume.

The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its keys must be defined
<b>secretName</b>	<b>string</b>	secretName is the name of the secret in the pod's namespace to use. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#secret">https://kubernetes.io/docs/concepts/storage/volumes#secret</a>

#### 14.13.1.308. .template.spec.volumes[].secret.items

##### Description

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys

will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.13.1.309. .template.spec.volumes[].secret.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

**14.13.1.310. .template.spec.volumes[].storageos****Description**

Represents a StorageOS persistent volume resource.

**Type****object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>volumeName</b>	<b>string</b>	volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
<b>volumeNamespace</b>	<b>string</b>	volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

#### 14.13.1.311. .template.spec.volumes[].storageos.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.13.1.312. .template.spec.volumes[].vsphereVolume

##### Description

Represents a vSphere volume resource.

##### Type

**object**

##### Required

- **volumePath**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>storagePolicyID</b>	<b>string</b>	storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
<b>storagePolicyName</b>	<b>string</b>	storagePolicyName is the storage Policy Based Management (SPBM) profile name.
<b>volumePath</b>	<b>string</b>	volumePath is the path that identifies vSphere volume vmdk

#### 14.13.2. API endpoints

The following API endpoints are available:

- **/api/v1/podtemplates**
  - **GET**: list or watch objects of kind PodTemplate
- **/api/v1/watch/podtemplates**

- **GET**: watch individual changes to a list of PodTemplate. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/podtemplates**
  - **DELETE**: delete collection of PodTemplate
  - **GET**: list or watch objects of kind PodTemplate
  - **POST**: create a PodTemplate
- **/api/v1/watch/namespaces/{namespace}/podtemplates**
  - **GET**: watch individual changes to a list of PodTemplate. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/podtemplates/{name}**
  - **DELETE**: delete a PodTemplate
  - **GET**: read the specified PodTemplate
  - **PATCH**: partially update the specified PodTemplate
  - **PUT**: replace the specified PodTemplate
- **/api/v1/watch/namespaces/{namespace}/podtemplates/{name}**
  - **GET**: watch changes to an object of kind PodTemplate. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 14.13.2.1. /api/v1/podtemplates

Table 14.342. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind PodTemplate

**Table 14.343. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>PodTemplateList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 14.13.2.2. /api/v1/watch/podtemplates

Table 14.344. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>"k8s.io/initial-events-end": "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PodTemplate. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.345. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.13.2.3. /api/v1/namespaces/{namespace}/podtemplates

Table 14.346. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.347. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of PodTemplate

Table 14.348. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.349. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.350. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind PodTemplate

Table 14.351. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.352. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PodTemplateList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a PodTemplate

**Table 14.353. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.354. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PodTemplate</b> schema	

**Table 14.355. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>PodTemplate</b> schema

HTTP code	Response body
201 - Created	<b>PodTemplate</b> schema
202 - Accepted	<b>PodTemplate</b> schema
401 - Unauthorized	Empty

#### 14.13.2.4. /api/v1/watch/namespaces/{namespace}/podtemplates

Table 14.356. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.357. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PodTemplate. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.358. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.13.2.5. /api/v1/namespaces/{namespace}/podtemplates/{name}

Table 14.359. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PodTemplate
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.360. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a PodTemplate

Table 14.361. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.362. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.363. HTTP responses

HTTP code	Response body
200 - OK	<b>PodTemplate</b> schema
202 - Accepted	<b>PodTemplate</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified PodTemplate

Table 14.364. HTTP responses

HTTP code	Response body
200 - OK	<b>PodTemplate</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified PodTemplate

**Table 14.365. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.366. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.367. HTTP responses

HTTP code	Response body
200 - OK	<b>PodTemplate</b> schema
201 - Created	<b>PodTemplate</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified PodTemplate

Table 14.368. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.369. Body parameters

Parameter	Type	Description
<b>body</b>	<b>PodTemplate</b> schema	

Table 14.370. HTTP responses

HTTP code	Response body
200 - OK	<b>PodTemplate</b> schema
201 - Created	<b>PodTemplate</b> schema
401 - Unauthorized	Empty

#### 14.13.2.6. /api/v1/watch/namespaces/{namespace}/podtemplates/{name}

Table 14.371. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PodTemplate
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.372. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind PodTemplate. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.373. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 14.14. REPLICATIONCONTROLLER [V1]

### Description

ReplicationController represents the configuration of a replication controller.

### Type

**object**

### 14.14.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	If the Labels of a ReplicationController are empty, they are defaulted to be the same as the Pod(s) that the replication controller manages. Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>spec</b>	<b>object</b>	ReplicationControllerSpec is the specification of a replication controller.
<b>status</b>	<b>object</b>	ReplicationControllerStatus represents the current status of a replication controller.

#### 14.14.1.1. .spec

##### Description

ReplicationControllerSpec is the specification of a replication controller.

##### Type

**object**

Property	Type	Description
<b>minReadySeconds</b>	<b>integer</b>	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
<b>replicas</b>	<b>integer</b>	Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#what-is-a-replicationcontroller">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#what-is-a-replicationcontroller</a>
<b>selector</b>	<b>object (string)</b>	Selector is a label query over pods that should match the Replicas count. If Selector is empty, it is defaulted to the labels present on the Pod template. Label keys and values that must match in order to be controlled by this replication controller, if empty defaulted to labels on Pod template. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors</a>

Property	Type	Description
<b>template</b>	<b>object</b>	PodTemplateSpec describes the data a pod should have when created from a template

#### 14.14.1.2. .spec.template

##### Description

PodTemplateSpec describes the data a pod should have when created from a template

##### Type

**object**

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PodSpec is a description of a pod.

#### 14.14.1.3. .spec.template.spec

##### Description

PodSpec is a description of a pod.

##### Type

**object**

##### Required

- **containers**

Property	Type	Description
<b>activeDeadlineSeconds</b>	<b>integer</b>	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
<b>affinity</b>	<b>object</b>	Affinity is a group of affinity scheduling rules.

Property	Type	Description
<b>automountServiceAccountToken</b>	<b>boolean</b>	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
<b>containers</b>	<b>array</b>	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
<b>containers[]</b>	<b>object</b>	A single application container that you want to run within a pod.
<b>dnsConfig</b>	<b>object</b>	PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

Property	Type	Description
<b>dnsPolicy</b>	<b>string</b>	<p>Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.</p> <p>Possible enum values: - <b>"ClusterFirst"</b> indicates that the pod should use cluster DNS first unless hostNetwork is true, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"ClusterFirstWithHostNet"</b> indicates that the pod should use cluster DNS first, if it is available, then fall back on the default (as determined by kubelet) DNS settings. - <b>"Default"</b> indicates that the pod should use the default (as determined by kubelet) DNS settings. - <b>"None"</b> indicates that the pod should use empty DNS settings. DNS parameters such as nameservers and search paths should be defined via DNSConfig.</p>
<b>enableServiceLinks</b>	<b>boolean</b>	<p>EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.</p>

Property	Type	Description
<b>ephemeralContainers</b>	<b>array</b>	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's <code>ephemeralcontainers</code> subresource.
<b>ephemeralContainers[]</b>	<b>object</b>	<p>An <code>EphemeralContainer</code> is a temporary container that you may add to an existing <code>Pod</code> for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a <code>Pod</code> is removed or restarted. The kubelet may evict a <code>Pod</code> if an ephemeral container causes the <code>Pod</code> to exceed its resource allocation.</p> <p>To add an ephemeral container, use the <code>ephemeralcontainers</code> subresource of an existing <code>Pod</code>. Ephemeral containers may not be removed or restarted.</p>
<b>hostAliases</b>	<b>array</b>	<code>HostAliases</code> is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non- <code>hostNetwork</code> pods.
<b>hostAliases[]</b>	<b>object</b>	<code>HostAlias</code> holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.
<b>hostIPC</b>	<b>boolean</b>	Use the host's <code>ipc</code> namespace. Optional: Default to false.



Property	Type	Description
<b>hostNetwork</b>	<b>boolean</b>	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
<b>hostPID</b>	<b>boolean</b>	Use the host's pid namespace. Optional: Default to false.
<b>hostUsers</b>	<b>boolean</b>	Use the host's user namespace. Optional: Default to true. If set to true or not present, the pod will be run in the host user namespace, useful for when the pod needs a feature only available to the host user namespace, such as loading a kernel module with CAP_SYS_MODULE. When set to false, a new users is created for the pod. Setting false is useful for mitigating container breakout vulnerabilities even allowing users to run their containers as root without actually having root privileges on the host. This field is alpha-level and is only honored by servers that enable the UserNamespacesSupport feature.
<b>hostname</b>	<b>string</b>	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
<b>imagePullSecrets</b>	<b>array</b>	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod">https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod</a>
<b>imagePullSecrets[]</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

Property	Type	Description
<b>initContainers</b>	<b>array</b>	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/init-containers/">https://kubernetes.io/docs/concepts/workloads/pods/init-containers/</a>
<b>initContainers[]</b>	<b>object</b>	A single application container that you want to run within a pod.
<b>nodeName</b>	<b>string</b>	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
<b>nodeSelector</b>	<b>object (string)</b>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a>

Property	Type	Description
<b>os</b>	<b>object</b>	PodOS defines the OS parameters of a pod.
<b>overhead</b>	<b>object (Quantity)</b>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md">https://git.k8s.io/enhancements/keps/sig-node/688-pod-overhead/README.md</a>
<b>preemptionPolicy</b>	<b>string</b>	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.  Possible enum values: - <b>"Never"</b> means that pod never preempts other pods with lower priority. - <b>"PreemptLowerPriority"</b> means that pod can preempt other pods with lower priority.
<b>priority</b>	<b>integer</b>	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

Property	Type	Description
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
<b>readinessGates</b>	<b>array</b>	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: <a href="https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates">https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates</a>
<b>readinessGates[]</b>	<b>object</b>	PodReadinessGate contains the reference to a pod condition
<b>resourceClaims</b>	<b>array</b>	ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable.
<b>resourceClaims[]</b>	<b>object</b>	PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

Property	Type	Description
<b>restartPolicy</b>	<b>string</b>	<p>Restart policy for all containers within the pod. One of Always, OnFailure, Never. In some contexts, only a subset of those values may be permitted. Default to Always. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy</a></p> <p>Possible enum values: - <b>"Always"</b> - <b>"Never"</b> - <b>"OnFailure"</b></p>
<b>runtimeClassName</b>	<b>string</b>	<p>RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: <a href="https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class">https://git.k8s.io/enhancements/keps/sig-node/585-runtime-class</a></p>
<b>schedulerName</b>	<b>string</b>	<p>If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.</p>

Property	Type	Description
<b>schedulingGates</b>	<b>array</b>	<p>SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.</p> <p>SchedulingGates can only be set at pod creation time, and be removed only afterwards.</p> <p>This is a beta feature enabled by the PodSchedulingReadiness feature gate.</p>
<b>schedulingGates[]</b>	<b>object</b>	PodSchedulingGate is associated to a Pod to guard its scheduling.
<b>securityContext</b>	<b>object</b>	PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.
<b>serviceAccount</b>	<b>string</b>	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
<b>serviceAccountName</b>	<b>string</b>	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: <a href="https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/">https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/</a>

Property	Type	Description
<b>setHostnameAsFQDN</b>	<b>boolean</b>	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
<b>shareProcessNamespace</b>	<b>boolean</b>	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
<b>subdomain</b>	<b>string</b>	If specified, the fully qualified Pod hostname will be "<hostname>.<subdomain>.<pod namespace>.svc.<cluster domain>". If not specified, the pod will not have a domainname at all.

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
<b>tolerations</b>	<b>array</b>	If specified, the pod's tolerations.
<b>tolerations[]</b>	<b>object</b>	The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.
<b>topologySpreadConstraints</b>	<b>array</b>	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
<b>topologySpreadConstraints[]</b>	<b>object</b>	TopologySpreadConstraint specifies how to spread matching pods among the given topology.
<b>volumes</b>	<b>array</b>	List of volumes that can be mounted by containers belonging to the pod. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes">https://kubernetes.io/docs/concepts/storage/volumes</a>
<b>volumes[]</b>	<b>object</b>	Volume represents a named volume in a pod that may be accessed by any container in the pod.



#### 14.14.1.4. .spec.template.spec.affinity

##### Description

Affinity is a group of affinity scheduling rules.

##### Type

**object**

Property	Type	Description
<b>nodeAffinity</b>	<b>object</b>	Node affinity is a group of node affinity scheduling rules.
<b>podAffinity</b>	<b>object</b>	Pod affinity is a group of inter pod affinity scheduling rules.
<b>podAntiAffinity</b>	<b>object</b>	Pod anti affinity is a group of inter pod anti affinity scheduling rules.

#### 14.14.1.5. .spec.template.spec.affinity.nodeAffinity

##### Description

Node affinity is a group of node affinity scheduling rules.

##### Type

**object**

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>object</b>	A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### 14.14.1.6. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, `requiredDuringScheduling` affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding `matchExpressions`; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.14.1.7. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

##### Description

An empty preferred scheduling term matches all objects with implicit weight 0 (i.e. it's a no-op). A null preferred scheduling term matches no objects (i.e. is also a no-op).

##### Type

**object**

##### Required

- **weight**
- **preference**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>preference</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.
<b>weight</b>	<b>integer</b>	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

#### 14.14.1.8. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.14.1.9. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

##### Description

A list of node selector requirements by node's labels.

##### Type

**array**

### 14.14.1.10. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.14.1.11. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

#### Description

A list of node selector requirements by node's fields.

#### Type

**array**

### 14.14.1.12. `.spec.template.spec.affinity.nodeAffinity.preferredDuringSchedulingIgnoredDuring`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

#### 14.14.1.13. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE

#### Description

A node selector represents the union of the results of one or more label queries over a set of nodes; that is, it represents the OR of the selectors represented by the node selector terms.

#### Type

**object**

#### Required

- **nodeSelectorTerms**

Property	Type	Description
<b>nodeSelectorTerms</b>	<b>array</b>	Required. A list of node selector terms. The terms are ORed.
<b>nodeSelectorTerms[]</b>	<b>object</b>	A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

#### 14.14.1.14. `.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE`

##### Description

Required. A list of node selector terms. The terms are ORed.

##### Type

**array**

#### 14.14.1.15. `.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE`

##### Description

A null or empty node selector term matches no objects. The requirements of them are ANDed. The TopologySelectorTerm type implements a subset of the NodeSelectorTerm.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of node selector requirements by node's labels.
<b>matchExpressions[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.
<b>matchFields</b>	<b>array</b>	A list of node selector requirements by node's fields.
<b>matchFields[]</b>	<b>object</b>	A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### 14.14.1.16. `.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE`

##### Description

A list of node selector requirements by node's labels.

#### Type

**array**

### 14.14.1.17. `.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE`

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.14.1.18. `.spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE`

#### Description

A list of node selector requirements by node's fields.

#### Type

**array**

### 14.14.1.19. .spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringE

#### Description

A node selector requirement is a selector that contains values, a key, and an operator that relates the key and values.

#### Type

**object**

#### Required

- **key**
- **operator**

Property	Type	Description
<b>key</b>	<b>string</b>	The label key that the selector applies to.
<b>operator</b>	<b>string</b>	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"Gt"</b> - <b>"In"</b> - <b>"Lt"</b> - <b>"NotIn"</b>
<b>values</b>	<b>array (string)</b>	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

### 14.14.1.20. .spec.template.spec.affinity.podAffinity

#### Description

Pod affinity is a group of inter pod affinity scheduling rules.

#### Type

**object**



Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

#### 14.14.1.21. .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringE

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.14.1.22. .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuringI

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

##### Type

**object**

##### Required

- **weight**
- **podAffinityTerm**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>podAffinityTerm</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

#### 14.14.1.23. .spec.template.spec.affinity.podAffinity.preferredDuringSchedulingIgnoredDuring

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.

Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.14.1.24. `.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringE`

##### Description

If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

##### Type

**array**

#### 14.14.1.25. `.spec.template.spec.affinity.podAffinity.requiredDuringSchedulingIgnoredDuringE`

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.14.1.26. .spec.template.spec.affinity.podAntiAffinity

##### Description

Pod anti affinity is a group of inter pod anti affinity scheduling rules.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>preferredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
<b>preferredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)
<b>requiredDuringSchedulingIgnoredDuringExecution</b>	<b>array</b>	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

Property	Type	Description
<b>requiredDuringSchedulingIgnoredDuringExecution[]</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

#### 14.14.1.27. .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDu

##### Description

The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

##### Type

**array**

#### 14.14.1.28. .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDu

##### Description

The weights of all of the matched WeightedPodAffinityTerm fields are added per-node to find the most preferred node(s)

##### Type

**object**

##### Required

- **weight**
- **podAffinityTerm**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>podAffinityTerm</b>	<b>object</b>	Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running
<b>weight</b>	<b>integer</b>	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

#### 14.14.1.29. .spec.template.spec.affinity.podAntiAffinity.preferredDuringSchedulingIgnoredDu

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.



Property	Type	Description
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.14.1.30. `.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDur`

##### Description

If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

##### Type

**array**

#### 14.14.1.31. `.spec.template.spec.affinity.podAntiAffinity.requiredDuringSchedulingIgnoredDurii`

##### Description

Defines a set of pods (namely those matching the labelSelector relative to the given namespace(s)) that this pod should be co-located (affinity) or not co-located (anti-affinity) with, where co-located is defined as running on a node whose value of the label with key <topologyKey> matches that of any node on which a pod of the set of pods is running

##### Type

**object**

##### Required

- **topologyKey**

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	A label query over a set of resources, in this case pods.
<b>namespaceSelector</b>	<b>LabelSelector</b>	A label query over the set of namespaces that the term applies to. The term is applied to the union of the namespaces selected by this field and the ones listed in the namespaces field. null selector and null or empty namespaces list means "this pod's namespace". An empty selector ({} ) matches all namespaces.
<b>namespaces</b>	<b>array (string)</b>	namespaces specifies a static list of namespace names that the term applies to. The term is applied to the union of the namespaces listed in this field and the ones selected by namespaceSelector. null or empty namespaces list and null namespaceSelector means "this pod's namespace".
<b>topologyKey</b>	<b>string</b>	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

#### 14.14.1.32. .spec.template.spec.containers

##### Description

List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.

##### Type

**array**

#### 14.14.1.33. .spec.template.spec.containers[]

##### Description

A single application container that you want to run within a pod.

Type

**object**

Required

- **name**

Property	Type	Description
<b>args</b>	<b>array (string)</b>	Arguments to the entrypoint. The container image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a>
<b>env</b>	<b>array</b>	List of environment variables to set in the container. Cannot be updated.
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps

Property	Type	Description
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
<b>imagePullPolicy</b>	<b>string</b>	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a>  Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present
<b>lifecycle</b>	<b>object</b>	Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.
<b>livenessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
<b>ports</b>	<b>array</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

Property	Type	Description
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.



Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.14.1.34. .spec.template.spec.containers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.14.1.35. .spec.template.spec.containers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.14.1.36. `..spec.template.spec.containers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources ( <code>cpu</code> , <code>memory</code> ) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.14.1.37. `..spec.template.spec.containers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a `ConfigMap`.

Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.14.1.38. `.spec.template.spec.containers[].env[].valueFrom.fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.14.1.39. `.spec.template.spec.containers[].env[].valueFrom.resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.14.1.40. .spec.template.spec.containers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.14.1.41. .spec.template.spec.containers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.14.1.42. .spec.template.spec.containers[].envFrom[]

##### Description

EnvFromSource represents the source of a set of ConfigMaps

### Type

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>object</b>	SecretEnvSource selects a Secret to populate the environment variables with.  The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### 14.14.1.43. .spec.template.spec.containers[].envFrom[].configMapRef

### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  
The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 14.14.1.44. .spec.template.spec.containers[].envFrom[].secretRef

## Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.14.1.45. .spec.template.spec.containers[].lifecycle

## Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

## Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.14.1.46. .spec.template.spec.containers[].lifecycle.postStart

## Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

**14.14.1.47. .spec.template.spec.containers[].lifecycle.postStart.exec****Description**

ExecAction describes a "run in container" action.

**Type**

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**14.14.1.48. .spec.template.spec.containers[].lifecycle.postStart.httpGet****Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

**object**

**Required**

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.14.1.49. .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.14.1.50. .spec.template.spec.containers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**



Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.51. `.spec.template.spec.containers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.52. `.spec.template.spec.containers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.14.1.53. `.spec.template.spec.containers[].lifecycle.preStop.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.54. `.spec.template.spec.containers[].lifecycle.preStop.httpGet`

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.14.1.55. `..spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders`

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.14.1.56. `..spec.template.spec.containers[].lifecycle.preStop.httpGet.httpHeaders[]`

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.57. .spec.template.spec.containers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.58. .spec.template.spec.containers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.59. .spec.template.spec.containers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.60. .spec.template.spec.containers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.61. .spec.template.spec.containers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.14.1.62. `.spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.14.1.63. `.spec.template.spec.containers[].livenessProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required



- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.64. `.spec.template.spec.containers[].livenessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.65. `.spec.template.spec.containers[].ports`

##### Description

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <https://github.com/kubernetes/kubernetes/issues/108255>. Cannot be updated.

##### Type

**array**

#### 14.14.1.66. `.spec.template.spec.containers[].ports[]`

**Description**

ContainerPort represents a network port in a single container.

**Type**

**object**

**Required**

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

**14.14.1.67. .spec.template.spec.containers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.68. .spec.template.spec.containers[].readinessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.69. .spec.template.spec.containers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.70. .spec.template.spec.containers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.14.1.71. `.spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.14.1.72. `.spec.template.spec.containers[].readinessProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.73. `.spec.template.spec.containers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.74. `.spec.template.spec.containers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.14.1.75. `.spec.template.spec.containers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.14.1.76. .spec.template.spec.containers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>



Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.14.1.77. `.spec.template.spec.containers[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.14.1.78. `.spec.template.spec.containers[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.14.1.79. `.spec.template.spec.containers[].securityContext`

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.14.1.80. .spec.template.spec.containers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities

Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.14.1.81. `.spec.template.spec.containers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.14.1.82. `.spec.template.spec.containers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.14.1.83. .spec.template.spec.containers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.14.1.84. .spec.template.spec.containers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket



Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.85. `..spec.template.spec.containers[].startupProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.86. .spec.template.spec.containers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.87. .spec.template.spec.containers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.14.1.88. `.spec.template.spec.containers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.14.1.89. `.spec.template.spec.containers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.90. .spec.template.spec.containers[].startupProbe.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.91. .spec.template.spec.containers[].volumeDevices

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.14.1.92. .spec.template.spec.containers[].volumeDevices[]

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.14.1.93. .spec.template.spec.containers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Cannot be updated.

## Type

**array**

**14.14.1.94. .spec.template.spec.containers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '!'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.14.1.95. .spec.template.spec.dnsConfig

##### Description

PodDNSConfig defines the DNS parameters of a pod in addition to those generated from DNSPolicy.

##### Type

**object**

Property	Type	Description
<b>nameservers</b>	<b>array (string)</b>	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
<b>options</b>	<b>array</b>	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
<b>options[]</b>	<b>object</b>	PodDNSConfigOption defines DNS resolver options of a pod.
<b>searches</b>	<b>array (string)</b>	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

#### 14.14.1.96. .spec.template.spec.dnsConfig.options

**Description**

A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

**Type**

**array**

**14.14.1.97. .spec.template.spec.dnsConfig.options[]****Description**

PodDNSConfigOption defines DNS resolver options of a pod.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Required.
<b>value</b>	<b>string</b>	

**14.14.1.98. .spec.template.spec.ephemeralContainers****Description**

List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource.

**Type**

**array**

**14.14.1.99. .spec.template.spec.ephemeralContainers[]****Description**

An EphemeralContainer is a temporary container that you may add to an existing Pod for user-initiated activities such as debugging. Ephemeral containers have no resource or scheduling guarantees, and they will not be restarted when they exit or when a Pod is removed or restarted. The kubelet may evict a Pod if an ephemeral container causes the Pod to exceed its resource allocation. To add an ephemeral container, use the ephemeralcontainers subresource of an existing Pod. Ephemeral containers may not be removed or restarted.

**Type**

**object**

**Required**

- **name**



Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Endpoint array. Not executed within a shell. The image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array</b>	<p>List of environment variables to set in the container. Cannot be updated.</p>

Property	Type	Description
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a>
<b>imagePullPolicy</b>	<b>string</b>	<p>Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a></p> <p>Possible enum values: -  <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present</p>

Property	Type	Description
<b>lifecycle</b>	<b>object</b>	Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.
<b>livenessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>name</b>	<b>string</b>	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
<b>ports</b>	<b>array</b>	Ports are not allowed for ephemeral containers.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.

Property	Type	Description
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

Property	Type	Description
<b>targetContainerName</b>	<b>string</b>	<p>If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container uses the namespaces configured in the Pod spec.</p> <p>The container runtime must implement support for this feature. If the runtime does not support namespace targeting then the result of setting this field is undefined.</p>
<b>terminationMessagePath</b>	<b>string</b>	<p>Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to <code>/dev/termination-log</code>. Cannot be updated.</p>

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.14.1.100. .spec.template.spec.ephemeralContainers[].env

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.14.1.101. .spec.template.spec.ephemeralContainers[].env[]

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.14.1.102. `.spec.template.spec.ephemeralContainers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources ( <code>cpu</code> , <code>memory</code> ) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.14.1.103. `.spec.template.spec.ephemeralContainers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a `ConfigMap`.



Type

**object**

Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.14.1.104. `.spec.template.spec.ephemeralContainers[].env[].valueFrom.fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.14.1.105. `.spec.template.spec.ephemeralContainers[].env[].valueFrom.resourceFieldRef`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.14.1.106. .spec.template.spec.ephemeralContainers[].env[].valueFrom.secretKeyRef

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.14.1.107. .spec.template.spec.ephemeralContainers[].envFrom

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.14.1.108. .spec.template.spec.ephemeralContainers[].envFrom[]

##### Description

EnvFromSource represents the source of a set of ConfigMaps

### Type

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.
<b>prefix</b>	<b>string</b>	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
<b>secretRef</b>	<b>object</b>	SecretEnvSource selects a Secret to populate the environment variables with.  The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### 14.14.1.109. .spec.template.spec.ephemeralContainers[].envFrom[].configMapRef

### Description

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.  
The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap must be defined

#### 14.14.1.110. .spec.template.spec.ephemeralContainers[].envFrom[].secretRef

## Description

SecretEnvSource selects a Secret to populate the environment variables with.

The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.14.1.111. .spec.template.spec.ephemeralContainers[].lifecycle

## Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

## Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.14.1.112. .spec.template.spec.ephemeralContainers[].lifecycle.postStart

## Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

## Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

**14.14.1.113. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.exec****Description**

ExecAction describes a "run in container" action.

**Type**

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

**14.14.1.114. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet****Description**

HTTPGetAction describes an action based on HTTP Get requests.

**Type**

**object**

**Required**

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.14.1.115. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeade

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.14.1.116. .spec.template.spec.ephemeralContainers[].lifecycle.postStart.httpGet.httpHeade

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.117. `.spec.template.spec.ephemeralContainers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.118. `.spec.template.spec.ephemeralContainers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPsocket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.14.1.119. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.120. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.



Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.14.1.121. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeader:

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.14.1.122. .spec.template.spec.ephemeralContainers[].lifecycle.preStop.httpGet.httpHeader

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.123. `.spec.template.spec.ephemeralContainers[].lifecycle.preStop.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.124. `.spec.template.spec.ephemeralContainers[].livenessProbe`

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.

Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.125. .spec.template.spec.ephemeralContainers[].livenessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.126. .spec.template.spec.ephemeralContainers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.127. .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.14.1.128. .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.14.1.129. .spec.template.spec.ephemeralContainers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.130. `.spec.template.spec.ephemeralContainers[].livenessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.131. `.spec.template.spec.ephemeralContainers[].ports`

##### Description

Ports are not allowed for ephemeral containers.

##### Type

**array**

#### 14.14.1.132. `.spec.template.spec.ephemeralContainers[].ports[]`

##### Description

ContainerPort represents a network port in a single container.

##### Type

**object**

## Required

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.

**14.14.1.133. .spec.template.spec.ephemeralContainers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type****object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.



Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.134. .spec.template.spec.ephemeralContainers[].readinessProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.135. .spec.template.spec.ephemeralContainers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.136. .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.14.1.137. .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.14.1.138. .spec.template.spec.ephemeralContainers[].readinessProbe.httpGet.httpHeaders

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.139. `.spec.template.spec.ephemeralContainers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.140. `.spec.template.spec.ephemeralContainers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.14.1.141. `.spec.template.spec.ephemeralContainers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

**14.14.1.142. .spec.template.spec.ephemeralContainers[].resources**

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.  This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.14.1.143. `.spec.template.spec.ephemeralContainers[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.14.1.144. `.spec.template.spec.ephemeralContainers[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.14.1.145. `.spec.template.spec.ephemeralContainers[].securityContext`

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

**Type****object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.



Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.14.1.146. .spec.template.spec.ephemeralContainers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities

Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.14.1.147. `.spec.template.spec.ephemeralContainers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.14.1.148. `.spec.template.spec.ephemeralContainers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.14.1.149. .spec.template.spec.ephemeralContainers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.14.1.150. .spec.template.spec.ephemeralContainers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.151. .spec.template.spec.ephemeralContainers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.152. .spec.template.spec.ephemeralContainers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.153. .spec.template.spec.ephemeralContainers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**



## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.14.1.154. `.spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.14.1.155. `.spec.template.spec.ephemeralContainers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.156. `.spec.template.spec.ephemeralContainers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.157. `.spec.template.spec.ephemeralContainers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.14.1.158. `.spec.template.spec.ephemeralContainers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type

**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.14.1.159. .spec.template.spec.ephemeralContainers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Subpath mounts are not allowed for ephemeral containers. Cannot be updated.

## Type

**array**

**14.14.1.160. .spec.template.spec.ephemeralContainers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '.'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: - <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). - <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.14.1.161. .spec.template.spec.hostAliases

##### Description

HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.

##### Type

**array**

#### 14.14.1.162. .spec.template.spec.hostAliases[]

##### Description

HostAlias holds the mapping between IP and hostnames that will be injected as an entry in the pod's hosts file.

##### Type

**object**

Property	Type	Description
<b>hostnames</b>	<b>array (string)</b>	Hostnames for the above IP address.
<b>ip</b>	<b>string</b>	IP address of the host file entry.

#### 14.14.1.163. .spec.template.spec.imagePullSecrets

##### Description

ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. More info:

<https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod>

##### Type

**array**

#### 14.14.1.164. .spec.template.spec.imagePullSecrets[]

## Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

### 14.14.1.165. .spec.template.spec.initContainers

## Description

List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: <https://kubernetes.io/docs/concepts/workloads/pods/init-containers/>

## Type

**array**

### 14.14.1.166. .spec.template.spec.initContainers[]

## Description

A single application container that you want to run within a pod.

## Type

**object**

## Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<p>Arguments to the endpoint. The container image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>command</b>	<b>array (string)</b>	<p>Entrypoint array. Not executed within a shell. The container image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. Double are reduced to a single \$, which allows for escaping the \$(VAR_NAME) syntax: i.e. "\$(VAR_NAME)" will produce the string literal "\$(VAR_NAME)". Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: <a href="https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell">https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell</a></p>
<b>env</b>	<b>array</b>	<p>List of environment variables to set in the container. Cannot be updated.</p>

Property	Type	Description
<b>env[]</b>	<b>object</b>	EnvVar represents an environment variable present in a Container.
<b>envFrom</b>	<b>array</b>	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
<b>envFrom[]</b>	<b>object</b>	EnvFromSource represents the source of a set of ConfigMaps
<b>image</b>	<b>string</b>	Container image name. More info: <a href="https://kubernetes.io/docs/concepts/containers/images">https://kubernetes.io/docs/concepts/containers/images</a> This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.



Property	Type	Description
<b>imagePullPolicy</b>	<b>string</b>	<p>Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: <a href="https://kubernetes.io/docs/concepts/containers/images#updating-images">https://kubernetes.io/docs/concepts/containers/images#updating-images</a></p> <p>Possible enum values: - <b>"Always"</b> means that kubelet always attempts to pull the latest image. Container will fail if the pull fails. - <b>"IfNotPresent"</b> means that kubelet pulls if the image isn't present on disk. Container will fail if the image isn't present and the pull fails. - <b>"Never"</b> means that kubelet never pulls an image, but only uses a local image. Container will fail if the image isn't present</p>
<b>lifecycle</b>	<b>object</b>	<p>Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.</p>
<b>livenessProbe</b>	<b>object</b>	<p>Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.</p>
<b>name</b>	<b>string</b>	<p>Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.</p>

Property	Type	Description
<b>ports</b>	<b>array</b>	List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <a href="https://github.com/kubernetes/kubernetes/issues/108255">https://github.com/kubernetes/kubernetes/issues/108255</a> . Cannot be updated.
<b>ports[]</b>	<b>object</b>	ContainerPort represents a network port in a single container.
<b>readinessProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.
<b>resizePolicy</b>	<b>array</b>	Resources resize policy for the container.
<b>resizePolicy[]</b>	<b>object</b>	ContainerResizePolicy represents resource resize policy for the container.
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>securityContext</b>	<b>object</b>	SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.
<b>startupProbe</b>	<b>object</b>	Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

Property	Type	Description
<b>stdin</b>	<b>boolean</b>	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
<b>stdinOnce</b>	<b>boolean</b>	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
<b>terminationMessagePath</b>	<b>string</b>	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

Property	Type	Description
<b>terminationMessagePolicy</b>	<b>string</b>	<p>Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.</p> <p>Possible enum values: - <b>"FallbackToLogsOnError"</b> will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents. - <b>"File"</b> is the default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits.</p>
<b>tty</b>	<b>boolean</b>	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
<b>volumeDevices</b>	<b>array</b>	volumeDevices is the list of block devices to be used by the container.
<b>volumeDevices[]</b>	<b>object</b>	volumeDevice describes a mapping of a raw block device within a container.
<b>volumeMounts</b>	<b>array</b>	Pod volumes to mount into the container's filesystem. Cannot be updated.

Property	Type	Description
<b>volumeMounts[]</b>	<b>object</b>	VolumeMount describes a mounting of a Volume within a container.
<b>workingDir</b>	<b>string</b>	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

#### 14.14.1.167. `.spec.template.spec.initContainers[].env`

##### Description

List of environment variables to set in the container. Cannot be updated.

##### Type

**array**

#### 14.14.1.168. `.spec.template.spec.initContainers[].env[]`

##### Description

EnvVar represents an environment variable present in a Container.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the environment variable. Must be a C_IDENTIFIER.

Property	Type	Description
<b>value</b>	<b>string</b>	Variable references <code>\$(VAR_NAME)</code> are expanded using the previously defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. Double <code>are</code> reduced to a single <code>\$</code> , which allows for escaping the <code>\$(VAR_NAME)</code> syntax: i.e. <code>"(VAR_NAME)"</code> will produce the string literal <code>"\$(VAR_NAME)"</code> . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to <code>""</code> .
<b>valueFrom</b>	<b>object</b>	<code>EnvVarSource</code> represents a source for the value of an <code>EnvVar</code> .

#### 14.14.1.169. `.spec.template.spec.initContainers[].env[].valueFrom`

##### Description

`EnvVarSource` represents a source for the value of an `EnvVar`.

##### Type

**object**

Property	Type	Description
<b>configMapKeyRef</b>	<b>object</b>	Selects a key from a <code>ConfigMap</code> .
<b>fieldRef</b>	<b>object</b>	<code>ObjectFieldSelector</code> selects an <code>APIVersioned</code> field of an object.
<b>resourceFieldRef</b>	<b>object</b>	<code>ResourceFieldSelector</code> represents container resources (cpu, memory) and their output format
<b>secretKeyRef</b>	<b>object</b>	<code>SecretKeySelector</code> selects a key of a <code>Secret</code> .

#### 14.14.1.170. `.spec.template.spec.initContainers[].env[].valueFrom.configMapKeyRef`

##### Description

Selects a key from a ConfigMap.

#### Type

**object**

#### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key to select.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the ConfigMap or its key must be defined

#### 14.14.1.171. .spec.template.spec.initContainers[].env[].valueFrom.fieldRef

#### Description

ObjectFieldSelector selects an APIVersioned field of an object.

#### Type

**object**

#### Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.14.1.172. .spec.template.spec.initContainers[].env[].valueFrom.resourceFieldRef

#### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### Type

**object**

#### Required

- **resource**

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.14.1.173. `.spec.template.spec.initContainers[].env[].valueFrom.secretKeyRef`

##### Description

SecretKeySelector selects a key of a Secret.

##### Type

**object**

##### Required

- **key**

Property	Type	Description
<b>key</b>	<b>string</b>	The key of the secret to select from. Must be a valid secret key.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret or its key must be defined

#### 14.14.1.174. `.spec.template.spec.initContainers[].envFrom`

##### Description

List of sources to populate environment variables in the container. The keys defined within a source must be a C\_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

##### Type

**array**

#### 14.14.1.175. `.spec.template.spec.initContainers[].envFrom[]`



**Description**

EnvFromSource represents the source of a set of ConfigMaps

**Type**

**object**

Property	Type	Description
<b>configMapRef</b>	<b>object</b>	<p>ConfigMapEnvSource selects a ConfigMap to populate the environment variables with.</p> <p>The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.</p>
<b>prefix</b>	<b>string</b>	<p>An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.</p>
<b>secretRef</b>	<b>object</b>	<p>SecretEnvSource selects a Secret to populate the environment variables with.</p> <p>The contents of the target Secret's Data field will represent the key-value pairs as environment variables.</p>

**14.14.1.176. .spec.template.spec.initContainers[].envFrom[].configMapRef****Description**

ConfigMapEnvSource selects a ConfigMap to populate the environment variables with. The contents of the target ConfigMap's Data field will represent the key-value pairs as environment variables.

**Type**

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	<p>Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a></p>
<b>optional</b>	<b>boolean</b>	<p>Specify whether the ConfigMap must be defined</p>

### 14.14.1.177. `.spec.template.spec.initContainers[].envFrom[].secretRef`

#### Description

SecretEnvSource selects a Secret to populate the environment variables with. The contents of the target Secret's Data field will represent the key-value pairs as environment variables.

#### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	Specify whether the Secret must be defined

### 14.14.1.178. `.spec.template.spec.initContainers[].lifecycle`

#### Description

Lifecycle describes actions that the management system should take in response to container lifecycle events. For the PostStart and PreStop lifecycle handlers, management of the container blocks until the action is complete, unless the container process fails, in which case the handler is aborted.

#### Type

**object**

Property	Type	Description
<b>postStart</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.
<b>preStop</b>	<b>object</b>	LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

### 14.14.1.179. `.spec.template.spec.initContainers[].lifecycle.postStart`

#### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCPSocket must be specified.

Type  
object

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.14.1.180. .spec.template.spec.initContainers[].lifecycle.postStart.exec

Description

ExecAction describes a "run in container" action.

Type  
object

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.181. .spec.template.spec.initContainers[].lifecycle.postStart.httpGet

Description

HTTPGetAction describes an action based on HTTP Get requests.

Type  
object

Required

- port

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.14.1.182. .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.14.1.183. .spec.template.spec.initContainers[].lifecycle.postStart.httpGet.httpHeaders[]

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.184. `.spec.template.spec.initContainers[].lifecycle.postStart.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.185. `.spec.template.spec.initContainers[].lifecycle.preStop`

##### Description

LifecycleHandler defines a specific action that should be taken in a lifecycle hook. One and only one of the fields, except TCP socket must be specified.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.

Property	Type	Description
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

#### 14.14.1.186. .spec.template.spec.initContainers[].lifecycle.preStop.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.187. .spec.template.spec.initContainers[].lifecycle.preStop.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.

Property	Type	Description
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

#### 14.14.1.188. `.spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders`

##### Description

Custom headers to set in the request. HTTP allows repeated headers.

##### Type

**array**

#### 14.14.1.189. `.spec.template.spec.initContainers[].lifecycle.preStop.httpGet.httpHeaders[]`

##### Description

HTTPHeader describes a custom header to be used in HTTP probes

##### Type

**object**

##### Required

- **name**
- **value**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.190. .spec.template.spec.initContainers[].lifecycle.preStop.tcpSocket

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.191. .spec.template.spec.initContainers[].livenessProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.



Property	Type	Description
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.192. `spec.template.spec.initContainers[].livenessProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.193. .spec.template.spec.initContainers[].livenessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.194. .spec.template.spec.initContainers[].livenessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.14.1.195. .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.14.1.196. .spec.template.spec.initContainers[].livenessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.197. `.spec.template.spec.initContainers[].livenessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.198. `.spec.template.spec.initContainers[].ports`

##### Description

List of ports to expose from the container. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Modifying this array with strategic merge patch may corrupt the data. For more information See <https://github.com/kubernetes/kubernetes/issues/108255>. Cannot be updated.

##### Type

**array**

#### 14.14.1.199. `.spec.template.spec.initContainers[].ports[]`

**Description**

ContainerPort represents a network port in a single container.

**Type**

**object**

**Required**

- **containerPort**

Property	Type	Description
<b>containerPort</b>	<b>integer</b>	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$ .
<b>hostIP</b>	<b>string</b>	What host IP to bind the external port to.
<b>hostPort</b>	<b>integer</b>	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$ . If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
<b>name</b>	<b>string</b>	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
<b>protocol</b>	<b>string</b>	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

**14.14.1.200. .spec.template.spec.initContainers[].readinessProbe****Description**

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

**Type**

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's <code>terminationGracePeriodSeconds</code> will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling <code>ProbeTerminationGracePeriod</code> feature gate. Minimum value is 1. <code>spec.terminationGracePeriodSeconds</code> is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.201. `.spec.template.spec.initContainers[].readinessProbe.exec`

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.202. .spec.template.spec.initContainers[].readinessProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.203. .spec.template.spec.initContainers[].readinessProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

## 14.14.1.204. .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

## 14.14.1.205. .spec.template.spec.initContainers[].readinessProbe.httpGet.httpHeaders[]

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.206. `.spec.template.spec.initContainers[].readinessProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.207. `.spec.template.spec.initContainers[].resizePolicy`

##### Description

Resources resize policy for the container.

##### Type

**array**

#### 14.14.1.208. `.spec.template.spec.initContainers[].resizePolicy[]`

##### Description

ContainerResizePolicy represents resource resize policy for the container.

##### Type

**object**

## Required

- **resourceName**
- **restartPolicy**

Property	Type	Description
<b>resourceName</b>	<b>string</b>	Name of the resource to which this resource resize policy applies. Supported values: cpu, memory.
<b>restartPolicy</b>	<b>string</b>	Restart policy to apply when specified resource is resized. If not specified, it defaults to NotRequired.

14.14.1.209. `.spec.template.spec.initContainers[].resources`

## Description

ResourceRequirements describes the compute resource requirements.

## Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	Claims lists the names of resources, defined in <code>spec.resourceClaims</code> , that are used by this container.  This is an alpha field and requires enabling the <code>DynamicResourceAllocation</code> feature gate.  This field is immutable. It can only be set for containers.
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in <code>PodSpec.ResourceClaims</code> .
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

Property	Type	Description
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.14.1.210. `.spec.template.spec.initContainers[].resources.claims`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.14.1.211. `.spec.template.spec.initContainers[].resources.claims[]`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in <code>pod.spec.resourceClaims</code> of the Pod where this field is used. It makes that resource available inside a container.

#### 14.14.1.212. `.spec.template.spec.initContainers[].securityContext`

##### Description

SecurityContext holds security configuration that will be applied to a container. Some fields are present in both SecurityContext and PodSecurityContext. When both are set, the values in SecurityContext take precedence.

## Type

**object**

Property	Type	Description
<b>allowPrivilegeEscalation</b>	<b>boolean</b>	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN Note that this field cannot be set when spec.os.name is windows.
<b>capabilities</b>	<b>object</b>	Adds and removes POSIX capabilities from running containers.
<b>privileged</b>	<b>boolean</b>	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>procMount</b>	<b>string</b>	<p>procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Default"</b> uses the container runtime defaults for readonly and masked paths for /proc. Most container runtimes mask certain paths in /proc to avoid accidental security exposure of special devices or information. -  <b>"Unmasked"</b> bypasses the default masking behavior of the container runtime and ensures the newly created /proc the container stays in tact with no modifications.</p>
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	Whether this container has a read-only root filesystem. Default is false. Note that this field cannot be set when spec.os.name is windows.
<b>runAsGroup</b>	<b>integer</b>	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.14.1.213. .spec.template.spec.initContainers[].securityContext.capabilities

##### Description

Adds and removes POSIX capabilities from running containers.

##### Type

**object**

Property	Type	Description
<b>add</b>	<b>array (string)</b>	Added capabilities



Property	Type	Description
<b>drop</b>	<b>array (string)</b>	Removed capabilities

#### 14.14.1.214. `.spec.template.spec.initContainers[].securityContext.seLinuxOptions`

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.14.1.215. `.spec.template.spec.initContainers[].securityContext.seccompProfile`

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.14.1.216. .spec.template.spec.initContainers[].securityContext.windowsOptions

##### Description

WindowsSecurityContextOptions contain Windows-specific options and credentials.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.14.1.217. .spec.template.spec.initContainers[].startupProbe

##### Description

Probe describes a health check to be performed against a container to determine whether it is alive or ready to receive traffic.

##### Type

**object**

Property	Type	Description
<b>exec</b>	<b>object</b>	ExecAction describes a "run in container" action.
<b>failureThreshold</b>	<b>integer</b>	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
<b>grpc</b>	<b>object</b>	GRPC specifies an action involving a GRPC port.
<b>httpGet</b>	<b>object</b>	HTTPGetAction describes an action based on HTTP Get requests.
<b>initialDelaySeconds</b>	<b>integer</b>	Number of seconds after the container has started before liveness probes are initiated. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>
<b>periodSeconds</b>	<b>integer</b>	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
<b>successThreshold</b>	<b>integer</b>	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
<b>tcpSocket</b>	<b>object</b>	TCPSocketAction describes an action based on opening a socket

Property	Type	Description
<b>terminationGracePeriodSeconds</b>	<b>integer</b>	Optional duration in seconds the pod needs to terminate gracefully upon probe failure. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. If this value is nil, the pod's terminationGracePeriodSeconds will be used. Otherwise, this value overrides the value provided by the pod spec. Value must be non-negative integer. The value zero indicates stop immediately via the kill signal (no opportunity to shut down). This is a beta field and requires enabling ProbeTerminationGracePeriod feature gate. Minimum value is 1. spec.terminationGracePeriodSeconds is used if unset.
<b>timeoutSeconds</b>	<b>integer</b>	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: <a href="https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes">https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes</a>

#### 14.14.1.218. .spec.template.spec.initContainers[].startupProbe.exec

##### Description

ExecAction describes a "run in container" action.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>command</b>	<b>array (string)</b>	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

#### 14.14.1.219. .spec.template.spec.initContainers[].startupProbe.grpc

##### Description

GRPC specifies an action involving a GRPC port.

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>port</b>	<b>integer</b>	Port number of the gRPC service. Number must be in the range 1 to 65535.
<b>service</b>	<b>string</b>	Service is the name of the service to place in the gRPC HealthCheckRequest (see <a href="https://github.com/grpc/grpc/blob/master/doc/health-checking.md">https://github.com/grpc/grpc/blob/master/doc/health-checking.md</a> ).  If this is not specified, the default behavior is defined by gRPC.

#### 14.14.1.220. .spec.template.spec.initContainers[].startupProbe.httpGet

##### Description

HTTPGetAction describes an action based on HTTP Get requests.

##### Type

**object**

## Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
<b>httpHeaders</b>	<b>array</b>	Custom headers to set in the request. HTTP allows repeated headers.
<b>httpHeaders[]</b>	<b>object</b>	HTTPHeader describes a custom header to be used in HTTP probes
<b>path</b>	<b>string</b>	Path to access on the HTTP server.
<b>port</b>	<b>IntOrString</b>	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
<b>scheme</b>	<b>string</b>	Scheme to use for connecting to the host. Defaults to HTTP.  Possible enum values: - <b>"HTTP"</b> means that the scheme used will be http:// - <b>"HTTPS"</b> means that the scheme used will be https://

14.14.1.221. `.spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders`

## Description

Custom headers to set in the request. HTTP allows repeated headers.

## Type

**array**

14.14.1.222. `.spec.template.spec.initContainers[].startupProbe.httpGet.httpHeaders[]`

## Description

HTTPHeader describes a custom header to be used in HTTP probes

## Type

**object**

## Required

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	The header field name. This will be canonicalized upon output, so case-variant names will be understood as the same header.
<b>value</b>	<b>string</b>	The header field value

#### 14.14.1.223. `.spec.template.spec.initContainers[].startupProbe.tcpSocket`

##### Description

TCPSocketAction describes an action based on opening a socket

##### Type

**object**

##### Required

- **port**

Property	Type	Description
<b>host</b>	<b>string</b>	Optional: Host name to connect to, defaults to the pod IP.
<b>port</b>	<b>IntOrString</b>	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

#### 14.14.1.224. `.spec.template.spec.initContainers[].volumeDevices`

##### Description

volumeDevices is the list of block devices to be used by the container.

##### Type

**array**

#### 14.14.1.225. `.spec.template.spec.initContainers[].volumeDevices[]`

##### Description

volumeDevice describes a mapping of a raw block device within a container.

##### Type



**object**

## Required

- **name**
- **devicePath**

Property	Type	Description
<b>devicePath</b>	<b>string</b>	devicePath is the path inside of the container that the device will be mapped to.
<b>name</b>	<b>string</b>	name must match the name of a persistentVolumeClaim in the pod

**14.14.1.226. .spec.template.spec.initContainers[].volumeMounts**

## Description

Pod volumes to mount into the container's filesystem. Cannot be updated.

## Type

**array**

**14.14.1.227. .spec.template.spec.initContainers[].volumeMounts[]**

## Description

VolumeMount describes a mounting of a Volume within a container.

## Type

**object**

## Required

- **name**
- **mountPath**

Property	Type	Description
<b>mountPath</b>	<b>string</b>	Path within the container at which the volume should be mounted. Must not contain '!'.

Property	Type	Description
<b>mountPropagation</b>	<b>string</b>	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.</p> <p>Possible enum values: -  <b>"Bidirectional"</b> means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology). -  <b>"HostToContainer"</b> means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology). - <b>"None"</b> means that the volume in a container will not receive new mounts from the host or other containers, and filesystems mounted inside the container won't be propagated to the host or other containers. Note that this mode corresponds to "private" in Linux terminology.</p>
<b>name</b>	<b>string</b>	This must match the Name of a Volume.
<b>readOnly</b>	<b>boolean</b>	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
<b>subPath</b>	<b>string</b>	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

Property	Type	Description
<b>subPathExpr</b>	<b>string</b>	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

#### 14.14.1.228. .spec.template.spec.os

##### Description

PodOS defines the OS parameters of a pod.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name is the name of the operating system. The currently supported values are linux and windows. Additional value may be defined in future and can be one of: <a href="https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration">https://github.com/opencontainers/runtime-spec/blob/master/config.md#platform-specific-configuration</a> Clients should expect to handle additional values and treat unrecognized values in this field as os: null

#### 14.14.1.229. .spec.template.spec.readinessGates

##### Description

If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True"  
More info: <https://git.k8s.io/enhancements/keps/sig-network/580-pod-readiness-gates>

##### Type

**array**

### 14.14.1.230. `.spec.template.spec.readinessGates[]`

#### Description

PodReadinessGate contains the reference to a pod condition

#### Type

**object**

#### Required

- **conditionType**

Property	Type	Description
<b>conditionType</b>	<b>string</b>	ConditionType refers to a condition in the pod's condition list with matching type.

### 14.14.1.231. `.spec.template.spec.resourceClaims`

#### Description

ResourceClaims defines which ResourceClaims must be allocated and reserved before the Pod is allowed to start. The resources will be made available to those containers which consume them by name.

This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.

This field is immutable.

#### Type

**array**

### 14.14.1.232. `.spec.template.spec.resourceClaims[]`

#### Description

PodResourceClaim references exactly one ResourceClaim through a ClaimSource. It adds a name to it that uniquely identifies the ResourceClaim inside the Pod. Containers that need access to the ResourceClaim reference it with this name.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name uniquely identifies this resource claim inside the pod. This must be a DNS_LABEL.

Property	Type	Description
<b>source</b>	<b>object</b>	<p>ClaimSource describes a reference to a ResourceClaim.</p> <p>Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.</p>

#### 14.14.1.233. .spec.template.spec.resourceClaims[].source

##### Description

ClaimSource describes a reference to a ResourceClaim.

Exactly one of these fields should be set. Consumers of this type must treat an empty object as if it has an unknown value.

##### Type

**object**

Property	Type	Description
<b>resourceClaimName</b>	<b>string</b>	ResourceClaimName is the name of a ResourceClaim object in the same namespace as this pod.

Property	Type	Description
<b>resourceClaimTemplateName</b>	<b>string</b>	<p>ResourceClaimTemplateName is the name of a ResourceClaimTemplate object in the same namespace as this pod.</p> <p>The template will be used to create a new ResourceClaim, which will be bound to this pod. When this pod is deleted, the ResourceClaim will also be deleted. The name of the ResourceClaim will be &lt;pod name&gt;-&lt;resource name&gt;, where &lt;resource name&gt; is the PodResourceClaim.Name. Pod validation will reject the pod if the concatenated name is not valid for a ResourceClaim (e.g. too long).</p> <p>An existing ResourceClaim with that name that is not owned by the pod will not be used for the pod to avoid using an unrelated resource by mistake. Scheduling and pod startup are then blocked until the unrelated ResourceClaim is removed.</p> <p>This field is immutable and no changes will be made to the corresponding ResourceClaim by the control plane after creating the ResourceClaim.</p>

#### 14.14.1.234. .spec.template.spec.schedulingGates

##### Description

SchedulingGates is an opaque list of values that if specified will block scheduling the pod. If schedulingGates is not empty, the pod will stay in the SchedulingGated state and the scheduler will not attempt to schedule the pod.

SchedulingGates can only be set at pod creation time, and be removed only afterwards.

This is a beta feature enabled by the PodSchedulingReadiness feature gate.

##### Type

**array**

#### 14.14.1.235. .spec.template.spec.schedulingGates[]

##### Description

PodSchedulingGate is associated to a Pod to guard its scheduling.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the scheduling gate. Each scheduling gate must have a unique name field.

### 14.14.1.236. .spec.template.spec.securityContext

#### Description

PodSecurityContext holds pod-level security attributes and common container settings. Some fields are also present in container.securityContext. Field values of container.securityContext take precedence over field values of PodSecurityContext.

#### Type

**object**

Property	Type	Description
<b>fsGroup</b>	<b>integer</b>	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none"> <li>1. The owning GID will be the FSGroup</li> <li>2. The setgid bit is set (new files created in the volume will be owned by FSGroup)</li> <li>3. The permission bits are OR'd with rw-rw----</li> </ol> <p>If unset, the Kubelet will not modify the ownership and permissions of any volume. Note that this field cannot be set when spec.os.name is windows.</p>

Property	Type	Description
<b>fsGroupChangePolicy</b>	<b>string</b>	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified, "Always" is used. Note that this field cannot be set when spec.os.name is windows.</p> <p>Possible enum values: -  <b>"Always"</b> indicates that volume's ownership and permissions should always be changed whenever volume is mounted inside a Pod. This the default behavior. -  <b>"OnRootMismatch"</b> indicates that volume's ownership and permissions will be changed only when permission and ownership of root directory does not match with expected permissions on the volume. This can help shorten the time it takes to change ownership and permissions of a volume.</p>
<b>runAsGroup</b>	<b>integer</b>	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.</p>



Property	Type	Description
<b>runAsNonRoot</b>	<b>boolean</b>	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
<b>runAsUser</b>	<b>integer</b>	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Note that this field cannot be set when spec.os.name is windows.
<b>seLinuxOptions</b>	<b>object</b>	SELinuxOptions are the labels to be applied to the container
<b>seccompProfile</b>	<b>object</b>	SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.
<b>supplementalGroups</b>	<b>array (integer)</b>	A list of groups applied to the first process run in each container, in addition to the container's primary GID, the fsGroup (if specified), and group memberships defined in the container image for the uid of the container process. If unspecified, no additional groups are added to any container. Note that group memberships defined in the container image for the uid of the container process are still effective, even if they are not included in this list. Note that this field cannot be set when spec.os.name is windows.

Property	Type	Description
<b>sysctls</b>	<b>array</b>	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.
<b>sysctls[]</b>	<b>object</b>	Sysctl defines a kernel parameter to be set
<b>windowsOptions</b>	<b>object</b>	WindowsSecurityContextOptions contain Windows-specific options and credentials.

#### 14.14.1.237. .spec.template.spec.securityContext.seLinuxOptions

##### Description

SELinuxOptions are the labels to be applied to the container

##### Type

**object**

Property	Type	Description
<b>level</b>	<b>string</b>	Level is SELinux level label that applies to the container.
<b>role</b>	<b>string</b>	Role is a SELinux role label that applies to the container.
<b>type</b>	<b>string</b>	Type is a SELinux type label that applies to the container.
<b>user</b>	<b>string</b>	User is a SELinux user label that applies to the container.

#### 14.14.1.238. .spec.template.spec.securityContext.seccompProfile

##### Description

SeccompProfile defines a pod/container's seccomp profile settings. Only one profile source may be set.

##### Type

**object**

##### Required

- **type**

Property	Type	Description
<b>localhostProfile</b>	<b>string</b>	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
<b>type</b>	<b>string</b>	<p>type indicates which kind of seccomp profile will be applied. Valid options are:</p> <p>Localhost - a profile defined in a file on the node should be used.  RuntimeDefault - the container runtime default profile should be used.  Unconfined - no profile should be applied.</p> <p>Possible enum values: -  <b>"Localhost"</b> indicates a profile defined in a file on the node should be used. The file's location relative to &lt;kubelet-root-dir&gt;/seccomp. -  <b>"RuntimeDefault"</b> represents the default container runtime seccomp profile. -  <b>"Unconfined"</b> indicates no seccomp profile is applied (A.K.A. unconfined).</p>

#### 14.14.1.239. .spec.template.spec.securityContext.sysctls

##### Description

Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch. Note that this field cannot be set when spec.os.name is windows.

##### Type

**array**

#### 14.14.1.240. .spec.template.spec.securityContext.sysctls[]

##### Description

Sysctl defines a kernel parameter to be set

##### Type

**object**

**Required**

- **name**
- **value**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of a property to set
<b>value</b>	<b>string</b>	Value of a property to set

**14.14.1.241. .spec.template.spec.securityContext.windowsOptions****Description**

WindowsSecurityContextOptions contain Windows-specific options and credentials.

**Type**

**object**

Property	Type	Description
<b>gmsaCredentialSpec</b>	<b>string</b>	GMSACredentialSpec is where the GMSA admission webhook ( <a href="https://github.com/kubernetes-sigs/windows-gmsa">https://github.com/kubernetes-sigs/windows-gmsa</a> ) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
<b>gmsaCredentialSpecName</b>	<b>string</b>	GMSACredentialSpecName is the name of the GMSA credential spec to use.

Property	Type	Description
<b>hostProcess</b>	<b>boolean</b>	HostProcess determines if a container should be run as a 'Host Process' container. This field is alpha-level and will only be honored by components that enable the WindowsHostProcessContainers feature flag. Setting this field without the feature flag will result in errors when validating the Pod. All of a Pod's containers must have the same effective HostProcess value (it is not allowed to have a mix of HostProcess containers and non-HostProcess containers). In addition, if HostProcess is true then HostNetwork must also be set to true.
<b>runAsUserName</b>	<b>string</b>	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

#### 14.14.1.242. .spec.template.spec.tolerations

##### Description

If specified, the pod's tolerations.

##### Type

**array**

#### 14.14.1.243. .spec.template.spec.tolerations[]

##### Description

The pod this Toleration is attached to tolerates any taint that matches the triple <key,value,effect> using the matching operator <operator>.

##### Type

**object**

Property	Type	Description
<b>effect</b>	<b>string</b>	<p>Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.</p> <p>Possible enum values: -  <b>"NoExecute"</b> Evict any already-running pods that do not tolerate the taint. Currently enforced by NodeController. -  <b>"NoSchedule"</b> Do not allow new pods to schedule onto the node unless they tolerate the taint, but allow all pods submitted to Kubelet without going through the scheduler to start, and allow all already-running pods to continue running. Enforced by the scheduler. -  <b>"PreferNoSchedule"</b> Like TaintEffectNoSchedule, but the scheduler tries not to schedule new pods onto the node, rather than prohibiting new pods from scheduling onto the node entirely. Enforced by the scheduler.</p>
<b>key</b>	<b>string</b>	<p>Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.</p>
<b>operator</b>	<b>string</b>	<p>Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.</p> <p>Possible enum values: - <b>"Equal"</b>  - <b>"Exists"</b></p>

Property	Type	Description
<b>tolerationSeconds</b>	<b>integer</b>	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
<b>value</b>	<b>string</b>	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

#### 14.14.1.244. .spec.template.spec.topologySpreadConstraints

##### Description

TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

##### Type

**array**

#### 14.14.1.245. .spec.template.spec.topologySpreadConstraints[]

##### Description

TopologySpreadConstraint specifies how to spread matching pods among the given topology.

##### Type

**object**

##### Required

- **maxSkew**
- **topologyKey**
- **whenUnsatisfiable**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>labelSelector</b>	<b>LabelSelector</b>	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.
<b>matchLabelKeys</b>	<b>array (string)</b>	<p>MatchLabelKeys is a set of pod label keys to select the pods over which spreading will be calculated. The keys are used to lookup values from the incoming pod labels, those key-value labels are ANDed with labelSelector to select the group of existing pods over which spreading will be calculated for the incoming pod. The same key is forbidden to exist in both MatchLabelKeys and LabelSelector. MatchLabelKeys cannot be set when LabelSelector isn't set. Keys that don't exist in the incoming pod labels will be ignored. A null or empty list means only match against labelSelector.</p> <p>This is a beta field and requires the MatchLabelKeysInPodTopologySpread feature gate to be enabled (enabled by default).</p>



Property	Type	Description
<b>maxSkew</b>	<b>integer</b>	<p>MaxSkew describes the degree to which pods may be unevenly distributed. When <b>whenUnsatisfiable=DoNotSchedule</b>, it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. The global minimum is the minimum number of matching pods in an eligible domain or zero if the number of eligible domains is less than MinDomains. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 2/2/1: In this case, the global minimum is 1.   zone1   zone2   zone3     P P   P P   P   - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 2/2/2; scheduling it onto zone1(zone2) would make the ActualSkew(3-1) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <b>whenUnsatisfiable=ScheduleAnyway</b>, it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.</p>

Property	Type	Description
<b>minDomains</b>	<b>integer</b>	<p>MinDomains indicates a minimum number of eligible domains. When the number of eligible domains with matching topology keys is less than minDomains, Pod Topology Spread treats "global minimum" as 0, and then the calculation of Skew is performed. And when the number of eligible domains with matching topology keys equals or greater than minDomains, this value has no effect on scheduling. As a result, when the number of eligible domains is less than minDomains, scheduler won't schedule more than maxSkew Pods to those domains. If value is nil, the constraint behaves as if MinDomains is equal to 1. Valid values are integers greater than 0. When value is not nil, WhenUnsatisfiable must be DoNotSchedule.</p> <p>For example, in a 3-zone cluster, MaxSkew is set to 2, MinDomains is set to 5 and pods with the same labelSelector spread as 2/2/2:   zone1   zone2   zone3     P P   P P   P P   The number of domains is less than 5(MinDomains), so "global minimum" is treated as 0. In this situation, new pod with the same labelSelector cannot be scheduled, because computed skew will be 3(3 - 0) if new Pod is scheduled to any of the three zones, it will violate MaxSkew.</p> <p>This is a beta field and requires the MinDomainsInPodTopologySpread feature gate to be enabled (enabled by default).</p>

Property	Type	Description
<b>nodeAffinityPolicy</b>	<b>string</b>	<p>NodeAffinityPolicy indicates how we will treat Pod's nodeAffinity/nodeSelector when calculating pod topology spread skew. Options are: - Honor: only nodes matching nodeAffinity/nodeSelector are included in the calculations. - Ignore: nodeAffinity/nodeSelector are ignored. All nodes are included in the calculations.</p> <p>If this value is nil, the behavior is equivalent to the Honor policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>
<b>nodeTaintsPolicy</b>	<b>string</b>	<p>NodeTaintsPolicy indicates how we will treat node taints when calculating pod topology spread skew. Options are: - Honor: nodes without taints, along with tainted nodes for which the incoming pod has a toleration, are included. - Ignore: node taints are ignored. All nodes are included.</p> <p>If this value is nil, the behavior is equivalent to the Ignore policy. This is a beta-level feature default enabled by the NodeInclusionPolicyInPodTopologySpread feature flag.</p> <p>Possible enum values: - <b>"Honor"</b> means use this scheduling directive when calculating pod topology spread skew. - <b>"Ignore"</b> means ignore this scheduling directive when calculating pod topology spread skew.</p>

Property	Type	Description
<b>topologyKey</b>	<b>string</b>	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. We define a domain as a particular instance of a topology. Also, we define an eligible domain as a domain whose nodes meet the requirements of nodeAffinityPolicy and nodeTaintsPolicy. e.g. If TopologyKey is "kubernetes.io/hostname", each Node is a domain of that topology. And, if TopologyKey is "topology.kubernetes.io/zone", each zone is a domain of that topology. It's a required field.

Property	Type	Description
<b>whenUnsatisfiable</b>	<b>string</b>	<p>WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1:   zone1   zone2   zone3     P P P   P   P   If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <b>more</b> imbalanced. It's a required field.</p> <p>Possible enum values: - <b>"DoNotSchedule"</b> instructs the scheduler not to schedule the pod when constraints are not satisfied. - <b>"ScheduleAnyway"</b> instructs the scheduler to schedule the pod even if constraints are not satisfied.</p>

#### 14.14.1.246. .spec.template.spec.volumes

##### Description

List of volumes that can be mounted by containers belonging to the pod. More info: <https://kubernetes.io/docs/concepts/storage/volumes>

##### Type

**array**

#### 14.14.1.247. .spec.template.spec.volumes[]

##### Description

Volume represents a named volume in a pod that may be accessed by any container in the pod.

## Type

**object**

## Required

- **name**

Property	Type	Description
<b>awsElasticBlockStore</b>	<b>object</b>	<p>Represents a Persistent Disk resource in AWS.</p> <p>An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.</p>
<b>azureDisk</b>	<b>object</b>	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
<b>azureFile</b>	<b>object</b>	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
<b>cephfs</b>	<b>object</b>	Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.
<b>cinder</b>	<b>object</b>	Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

Property	Type	Description
<b>configMap</b>	<b>object</b>	<p>Adapts a ConfigMap into a volume.</p> <p>The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.</p>
<b>csi</b>	<b>object</b>	Represents a source location of a volume to mount, managed by an external CSI driver
<b>downwardAPI</b>	<b>object</b>	DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.
<b>emptyDir</b>	<b>object</b>	Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.
<b>ephemeral</b>	<b>object</b>	Represents an ephemeral volume that is handled by a normal storage driver.
<b>fc</b>	<b>object</b>	Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.
<b>flexVolume</b>	<b>object</b>	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

Property	Type	Description
<b>flocker</b>	<b>object</b>	Represents a Flocker volume mounted by the Flocker agent. One and only one of datasetName and datasetUUID should be set. Flocker volumes do not support ownership management or SELinux relabeling.
<b>gcePersistentDisk</b>	<b>object</b>	Represents a Persistent Disk resource in Google Compute Engine.  A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.
<b>gitRepo</b>	<b>object</b>	Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.  DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
<b>glusterfs</b>	<b>object</b>	Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.
<b>hostPath</b>	<b>object</b>	Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.



Property	Type	Description
<b>iscsi</b>	<b>object</b>	Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.
<b>name</b>	<b>string</b>	name of the volume. Must be a DNS_LABEL and unique within the pod. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>nfs</b>	<b>object</b>	Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.
<b>persistentVolumeClaim</b>	<b>object</b>	PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).
<b>photonPersistentDisk</b>	<b>object</b>	Represents a Photon Controller persistent disk resource.
<b>portworxVolume</b>	<b>object</b>	PortworxVolumeSource represents a Portworx volume resource.
<b>projected</b>	<b>object</b>	Represents a projected volume source
<b>quobyte</b>	<b>object</b>	Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

Property	Type	Description
<b>rbd</b>	<b>object</b>	Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.
<b>scaleIO</b>	<b>object</b>	ScaleIOVolumeSource represents a persistent ScaleIO volume
<b>secret</b>	<b>object</b>	Adapts a Secret into a volume.  The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.
<b>storageos</b>	<b>object</b>	Represents a StorageOS persistent volume resource.
<b>vsphereVolume</b>	<b>object</b>	Represents a vSphere volume resource.

#### 14.14.1.248. .spec.template.spec.volumes[].awsElasticBlockStore

##### Description

Represents a Persistent Disk resource in AWS.

An AWS EBS disk must exist before mounting to a container. The disk must also be in the same AWS zone as the kubelet. An AWS EBS disk can only be mounted as read/write once. AWS EBS volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumelD**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
<b>readOnly</b>	<b>boolean</b>	readOnly value true will force the readOnly setting in VolumeMounts. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>
<b>volumeID</b>	<b>string</b>	volumeID is unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore">https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore</a>

#### 14.14.1.249. .spec.template.spec.volumes[].azureDisk

##### Description

AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **diskName**
- **diskURI**

Property	Type	Description
<b>cachingMode</b>	<b>string</b>	cachingMode is the Host Caching mode: None, Read Only, Read Write.  Possible enum values: - <b>"None"</b> - <b>"ReadOnly"</b> - <b>"ReadWrite"</b>
<b>diskName</b>	<b>string</b>	diskName is the Name of the data disk in the blob storage
<b>diskURI</b>	<b>string</b>	diskURI is the URI of data disk in the blob storage
<b>fsType</b>	<b>string</b>	fsType is Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>kind</b>	<b>string</b>	kind expected values are Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared  Possible enum values: - <b>"Dedicated"</b> - <b>"Managed"</b> - <b>"Shared"</b>
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

#### 14.14.1.250. .spec.template.spec.volumes[].azureFile

##### Description

AzureFile represents an Azure File Service mount on the host and bind mount to the pod.

##### Type

**object**

##### Required

- **secretName**
- **shareName**

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretName</b>	<b>string</b>	secretName is the name of secret that contains Azure Storage Account Name and Key
<b>shareName</b>	<b>string</b>	shareName is the azure share Name

#### 14.14.1.251. .spec.template.spec.volumes[].cephfs

##### Description

Represents a Ceph Filesystem mount that lasts the lifetime of a pod Cephfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**

Property	Type	Description
<b>monitors</b>	<b>array (string)</b>	monitors is Required: Monitors is a collection of Ceph monitors More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>path</b>	<b>string</b>	path is Optional: Used as the mounted root, rather than the full Ceph tree, default is /
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

Property	Type	Description
<b>secretFile</b>	<b>string</b>	secretFile is Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>user</b>	<b>string</b>	user is optional: User is the rados user name, default is admin More info: <a href="https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it">https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it</a>

#### 14.14.1.252. .spec.template.spec.volumes[].cephfs.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.14.1.253. .spec.template.spec.volumes[].cinder

##### Description

Represents a cinder volume resource in Openstack. A Cinder volume must exist before mounting to a container. The volume must also be in the same region as the kubelet. Cinder volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **volumeID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>volumeID</b>	<b>string</b>	volumeID used to identify the volume in cinder. More info: <a href="https://examples.k8s.io/mysql-cinder-pd/README.md">https://examples.k8s.io/mysql-cinder-pd/README.md</a>

#### 14.14.1.254. .spec.template.spec.volumes[].cinder.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.14.1.255. .spec.template.spec.volumes[].configMap

##### Description

Adapts a ConfigMap into a volume.

The contents of the target ConfigMap's Data field will be presented in a volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. ConfigMap volumes support ownership management and SELinux relabeling.

Type  
object

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

#### 14.14.1.256. .spec.template.spec.volumes[].configMap.items

##### Description



items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.14.1.257. .spec.template.spec.volumes[].configMap.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

**14.14.1.258. .spec.template.spec.volumes[].csi****Description**

Represents a source location of a volume to mount, managed by an external CSI driver

**Type**

**object**

## Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
<b>fsType</b>	<b>string</b>	fsType to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
<b>nodePublishSecretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>readOnly</b>	<b>boolean</b>	readOnly specifies a read-only configuration for the volume. Defaults to false (read/write).
<b>volumeAttributes</b>	<b>object (string)</b>	volumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

**14.14.1.259. .spec.template.spec.volumes[].csi.nodePublishSecretRef**

## Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

## Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

14.14.1.260. `.spec.template.spec.volumes[].downwardAPI`**Description**

DownwardAPIVolumeSource represents a volume containing downward API info. Downward API volumes support ownership management and SELinux relabeling.

**Type**

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	Items is a list of downward API volume file
<b>items[]</b>	<b>object</b>	DownwardAPIVolumeFile represents information to create the file containing the pod field

14.14.1.261. `.spec.template.spec.volumes[].downwardAPI.items`**Description**

Items is a list of downward API volume file

**Type**

**array**

14.14.1.262. `.spec.template.spec.volumes[].downwardAPI.items[]`**Description**

DownwardAPIVolumeFile represents information to create the file containing the pod field

**Type**

**object**

**Required**

**path**

- **path**

Property	Type	Description
<b>fieldRef</b>	<b>object</b>	ObjectFieldSelector selects an APIVersioned field of an object.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
<b>resourceFieldRef</b>	<b>object</b>	ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### 14.14.1.263. .spec.template.spec.volumes[].downwardAPI.items[].fieldRef

##### Description

ObjectFieldSelector selects an APIVersioned field of an object.

##### Type

**object**

##### Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.14.1.264. `.spec.template.spec.volumes[].downwardAPI.items[].resourceFieldRef`

##### Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

##### Type

**object**

##### Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.14.1.265. `.spec.template.spec.volumes[].emptyDir`

##### Description

Represents an empty directory for a pod. Empty directory volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>medium</b>	<b>string</b>	medium represents what type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

Property	Type	Description
<b>sizeLimit</b>	<b>Quantity</b>	sizeLimit is the total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#emptydir">https://kubernetes.io/docs/concepts/storage/volumes#emptydir</a>

#### 14.14.1.266. .spec.template.spec.volumes[].ephemeral

##### Description

Represents an ephemeral volume that is handled by a normal storage driver.

##### Type

**object**

Property	Type	Description
<b>volumeClaimTemplate</b>	<b>object</b>	PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

#### 14.14.1.267. .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate

##### Description

PersistentVolumeClaimTemplate is used to produce PersistentVolumeClaim objects as part of an EphemeralVolumeSource.

##### Type

**object**

##### Required

- **spec**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
<b>spec</b>	<b>object</b>	PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

#### 14.14.1.268. .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec

##### Description

PersistentVolumeClaimSpec describes the common attributes of storage devices and allows a Source for provider-specific attributes

##### Type

**object**

Property	Type	Description
<b>accessModes</b>	<b>array (string)</b>	accessModes contains the desired access modes the volume should have. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1</a>
<b>dataSource</b>	<b>object</b>	TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

Property	Type	Description
<b>dataSourceRef</b>	<b>object</b>	<p>dataSourceRef specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a PersistentVolumeClaim object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the dataSource field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in dataSourceRef, both fields (dataSource and dataSourceRef) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in dataSourceRef, dataSource isn't set to the same value and must be empty. There are three important differences between dataSource and dataSourceRef: * While dataSource only allows two specific types of objects, dataSourceRef allows any non-core object, as well as PersistentVolumeClaim objects. * While dataSource ignores disallowed values (dropping them), dataSourceRef preserves all values, and generates an error if a disallowed value is specified. * While dataSource only allows local objects, dataSourceRef allows objects in any namespaces. (Beta) Using this field requires the AnyVolumeDataSource feature gate to be enabled. (Alpha) Using the namespace field of dataSourceRef requires the CrossNamespaceVolumeDataSource feature gate to be enabled.</p>



Property	Type	Description
<b>resources</b>	<b>object</b>	ResourceRequirements describes the compute resource requirements.
<b>selector</b>	<b>LabelSelector</b>	selector is a label query over volumes to consider for binding.
<b>storageClassName</b>	<b>string</b>	storageClassName is the name of the StorageClass required by the claim. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1">https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1</a>
<b>volumeMode</b>	<b>string</b>	<p>volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.</p> <p>Possible enum values: - <b>"Block"</b> means the volume will not be formatted with a filesystem and will remain a raw block device. - <b>"Filesystem"</b> means the volume will be or is formatted with a filesystem.</p>
<b>volumeName</b>	<b>string</b>	volumeName is the binding reference to the PersistentVolume backing this claim.

#### 14.14.1.269. .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSource

##### Description

TypedLocalObjectReference contains enough information to let you locate the typed referenced object inside the same namespace.

##### Type

**object**

##### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

#### 14.14.1.270. `.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.dataSourceRef`

##### Description

`dataSourceRef` specifies the object from which to populate the volume with data, if a non-empty volume is desired. This may be any object from a non-empty API group (non core object) or a `PersistentVolumeClaim` object. When this field is specified, volume binding will only succeed if the type of the specified object matches some installed volume populator or dynamic provisioner. This field will replace the functionality of the `dataSource` field and as such if both fields are non-empty, they must have the same value. For backwards compatibility, when namespace isn't specified in `dataSourceRef`, both fields (`dataSource` and `dataSourceRef`) will be set to the same value automatically if one of them is empty and the other is non-empty. When namespace is specified in `dataSourceRef`, `dataSource` isn't set to the same value and must be empty. There are three important differences between `dataSource` and `dataSourceRef`: \* While `dataSource` only allows two specific types of objects, `dataSourceRef` allows any non-core object, as well as `PersistentVolumeClaim` objects. \* While `dataSource` ignores disallowed values (dropping them), `dataSourceRef` preserves all values, and generates an error if a disallowed value is specified. \* While `dataSource` only allows local objects, `dataSourceRef` allows objects in any namespaces. (Beta) Using this field requires the `AnyVolumeDataSource` feature gate to be enabled. (Alpha) Using the namespace field of `dataSourceRef` requires the `CrossNamespaceVolumeDataSource` feature gate to be enabled.

##### Type

**object**

##### Required

- **kind**
- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced
<b>namespace</b>	<b>string</b>	Namespace is the namespace of resource being referenced Note that when a namespace is specified, a gateway.networking.k8s.io/ReferenceGrant object is required in the referent namespace to allow that namespace's owner to accept the reference. See the ReferenceGrant documentation for details. (Alpha) This field requires the CrossNamespaceVolumeDataSource feature gate to be enabled.

#### 14.14.1.271. .spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources

##### Description

ResourceRequirements describes the compute resource requirements.

##### Type

**object**

Property	Type	Description
<b>claims</b>	<b>array</b>	<p>Claims lists the names of resources, defined in spec.resourceClaims, that are used by this container.</p> <p>This is an alpha field and requires enabling the DynamicResourceAllocation feature gate.</p> <p>This field is immutable. It can only be set for containers.</p>

Property	Type	Description
<b>claims[]</b>	<b>object</b>	ResourceClaim references one entry in PodSpec.ResourceClaims.
<b>limits</b>	<b>object (Quantity)</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>object (Quantity)</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

#### 14.14.1.272. `.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.c`

##### Description

Claims lists the names of resources, defined in `spec.resourceClaims`, that are used by this container. This is an alpha field and requires enabling the `DynamicResourceAllocation` feature gate.

This field is immutable. It can only be set for containers.

##### Type

**array**

#### 14.14.1.273. `.spec.template.spec.volumes[].ephemeral.volumeClaimTemplate.spec.resources.c`

##### Description

ResourceClaim references one entry in `PodSpec.ResourceClaims`.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	Name must match the name of one entry in pod.spec.resourceClaims of the Pod where this field is used. It makes that resource available inside a container.

#### 14.14.1.274. .spec.template.spec.volumes[].fc

##### Description

Represents a Fibre Channel volume. Fibre Channel volumes can only be mounted as read/write once. Fibre Channel volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>lun</b>	<b>integer</b>	lun is Optional: FC target lun number
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>targetWWNs</b>	<b>array (string)</b>	targetWWNs is Optional: FC target worldwide names (WWNs)
<b>wwids</b>	<b>array (string)</b>	wwids Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

#### 14.14.1.275. .spec.template.spec.volumes[].flexVolume

##### Description

FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.

##### Type

**object**

## Required

- **driver**

Property	Type	Description
<b>driver</b>	<b>string</b>	driver is the name of the driver to use for this volume.
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
<b>options</b>	<b>object (string)</b>	options is Optional: this field holds extra command options if any.
<b>readOnly</b>	<b>boolean</b>	readOnly is Optional: defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**14.14.1.276. .spec.template.spec.volumes[].flexVolume.secretRef****Description**

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

**Type****object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

**14.14.1.277. .spec.template.spec.volumes[].flocker****Description**

Represents a Flocker volume mounted by the Flocker agent. One and only one of `datasetName` and `datasetUUID` should be set. Flocker volumes do not support ownership management or SELinux relabeling.

### Type

**object**

Property	Type	Description
<b>datasetName</b>	<b>string</b>	<code>datasetName</code> is Name of the dataset stored as metadata → name on the dataset for Flocker should be considered as deprecated
<b>datasetUUID</b>	<b>string</b>	<code>datasetUUID</code> is the UUID of the dataset. This is unique identifier of a Flocker dataset

### 14.14.1.278. `.spec.template.spec.volumes[].gcePersistentDisk`

#### Description

Represents a Persistent Disk resource in Google Compute Engine.

A GCE PD must exist before mounting to a container. The disk must also be in the same GCE project and zone as the kubelet. A GCE PD can only be mounted as read/write once or read-only many times. GCE PDs support ownership management and SELinux relabeling.

### Type

**object**

#### Required

- `pdName`

Property	Type	Description
<b>fsType</b>	<b>string</b>	<code>fsType</code> is filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

Property	Type	Description
<b>partition</b>	<b>integer</b>	partition is the partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>pdName</b>	<b>string</b>	pdName is unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk">https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</a>

#### 14.14.1.279. .spec.template.spec.volumes[].gitRepo

##### Description

Represents a volume that is populated with the contents of a git repository. Git repo volumes do not support ownership management. Git repo volumes support SELinux relabeling.

DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.

##### Type

**object**

##### Required

- **repository**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>directory</b>	<b>string</b>	directory is the target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
<b>repository</b>	<b>string</b>	repository is the URL
<b>revision</b>	<b>string</b>	revision is the commit hash for the specified revision.

#### 14.14.1.280. .spec.template.spec.volumes[].glusterfs

##### Description

Represents a Glusterfs mount that lasts the lifetime of a pod. Glusterfs volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **endpoints**
- **path**

Property	Type	Description
<b>endpoints</b>	<b>string</b>	endpoints is the endpoint name that details Glusterfs topology. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>
<b>path</b>	<b>string</b>	path is the Glusterfs volume path. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod">https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod</a>

#### 14.14.1.281. .spec.template.spec.volumes[].hostPath

##### Description

Represents a host path mapped into a pod. Host path volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a>

Property	Type	Description
<b>type</b>	<b>string</b>	<p>type for HostPath Volume Defaults to "" More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#hostpath">https://kubernetes.io/docs/concepts/storage/volumes#hostpath</a></p> <p>Possible enum values: - "" For backwards compatible, leave it empty if unset - <b>"BlockDevice"</b> A block device must exist at the given path - <b>"CharDevice"</b> A character device must exist at the given path - <b>"Directory"</b> A directory must exist at the given path - <b>"DirectoryOrCreate"</b> If nothing exists at the given path, an empty directory will be created there as needed with file mode 0755, having the same group and ownership with Kubelet. - <b>"File"</b> A file must exist at the given path - <b>"FileOrCreate"</b> If nothing exists at the given path, an empty file will be created there as needed with file mode 0644, having the same group and ownership with Kubelet. - <b>"Socket"</b> A UNIX socket must exist at the given path</p>

#### 14.14.1.282. .spec.template.spec.volumes[].iscsi

##### Description

Represents an iSCSI disk. iSCSI volumes can only be mounted as read/write once. iSCSI volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **targetPortal**
- **iqn**
- **lun**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>chapAuthDiscovery</b>	<b>boolean</b>	chapAuthDiscovery defines whether support iSCSI Discovery CHAP authentication
<b>chapAuthSession</b>	<b>boolean</b>	chapAuthSession defines whether support iSCSI Session CHAP authentication
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#iscsi">https://kubernetes.io/docs/concepts/storage/volumes#iscsi</a>
<b>initiatorName</b>	<b>string</b>	initiatorName is the custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <target portal>:<volume name> will be created for the connection.
<b>iqn</b>	<b>string</b>	iqn is the target iSCSI Qualified Name.
<b>iscsilInterface</b>	<b>string</b>	iscsilInterface is the interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
<b>lun</b>	<b>integer</b>	lun represents iSCSI Target Lun number.
<b>portals</b>	<b>array (string)</b>	portals is the iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

Property	Type	Description
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>targetPortal</b>	<b>string</b>	targetPortal is iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

#### 14.14.1.283. .spec.template.spec.volumes[].iscsi.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.14.1.284. .spec.template.spec.volumes[].nfs

##### Description

Represents an NFS mount that lasts the lifetime of a pod. NFS volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **server**
- **path**

Property	Type	Description
<b>path</b>	<b>string</b>	path that is exported by the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

Property	Type	Description
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>
<b>server</b>	<b>string</b>	server is the hostname or IP address of the NFS server. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#nfs">https://kubernetes.io/docs/concepts/storage/volumes#nfs</a>

#### 14.14.1.285. .spec.template.spec.volumes[].persistentVolumeClaim

##### Description

PersistentVolumeClaimVolumeSource references the user's PVC in the same namespace. This volume finds the bound PV and mounts that volume for the pod. A

PersistentVolumeClaimVolumeSource is, essentially, a wrapper around another type of volume that is owned by someone else (the system).

##### Type

**object**

##### Required

- **claimName**

Property	Type	Description
<b>claimName</b>	<b>string</b>	claimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims">https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims</a>
<b>readOnly</b>	<b>boolean</b>	readOnly Will force the ReadOnly setting in VolumeMounts. Default false.

#### 14.14.1.286. .spec.template.spec.volumes[].photonPersistentDisk

##### Description

Represents a Photon Controller persistent disk resource.

##### Type

**object**

**Required**

- **pdID**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>pdID</b>	<b>string</b>	pdID is the ID that identifies Photon Controller persistent disk

**14.14.1.287. .spec.template.spec.volumes[].portworxVolume****Description**

PortworxVolumeSource represents a Portworx volume resource.

**Type**

**object**

**Required**

- **volumelD**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>volumelD</b>	<b>string</b>	volumelD uniquely identifies a Portworx volume

**14.14.1.288. .spec.template.spec.volumes[].projected****Description**

Represents a projected volume source

## Type

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode are the mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>sources</b>	<b>array</b>	sources is the list of volume projections
<b>sources[]</b>	<b>object</b>	Projection that may be projected along with other supported volume types

14.14.1.289. `.spec.template.spec.volumes[].projected.sources`

## Description

sources is the list of volume projections

## Type

**array**14.14.1.290. `.spec.template.spec.volumes[].projected.sources[]`

## Description

Projection that may be projected along with other supported volume types

## Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>configMap</b>	<b>object</b>	<p>Adapts a ConfigMap into a projected volume.</p> <p>The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.</p>
<b>downwardAPI</b>	<b>object</b>	<p>Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.</p>
<b>secret</b>	<b>object</b>	<p>Adapts a secret into a projected volume.</p> <p>The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.</p>
<b>serviceAccountToken</b>	<b>object</b>	<p>ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).</p>

#### 14.14.1.291. .spec.template.spec.volumes[].projected.sources[].configMap

##### Description

Adapts a ConfigMap into a projected volume.

The contents of the target ConfigMap's Data field will be presented in a projected volume as files using the keys in the Data field as the file names, unless the items element is populated with specific mappings of keys to paths. Note that this is identical to a configmap volume source without the default mode.

##### Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional specify whether the ConfigMap or its keys must be defined

**14.14.1.292. .spec.template.spec.volumes[].projected.sources[].configMap.items****Description**

items if unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.14.1.293. .spec.template.spec.volumes[].projected.sources[].configMap.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.14.1.294. `.spec.template.spec.volumes[].projected.sources[].downwardAPI`

##### Description

Represents downward API info for projecting into a projected volume. Note that this is identical to a downwardAPI volume source without the default mode.

##### Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	Items is a list of DownwardAPIVolume file
<b>items[]</b>	<b>object</b>	DownwardAPIVolumeFile represents information to create the file containing the pod field

#### 14.14.1.295. `.spec.template.spec.volumes[].projected.sources[].downwardAPI.items`

##### Description

Items is a list of DownwardAPIVolume file

Type

**array**

#### 14.14.1.296. `.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[]`

Description

DownwardAPIVolumeFile represents information to create the file containing the pod field

Type

**object**

Required

- **path**

Property	Type	Description
<b>fieldRef</b>	<b>object</b>	ObjectFieldSelector selects an APIVersioned field of an object.
<b>mode</b>	<b>integer</b>	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the <code>..</code> path. Must be utf-8 encoded. The first item of the relative path must not start with <code>..</code>
<b>resourceFieldRef</b>	<b>object</b>	ResourceFieldSelector represents container resources (cpu, memory) and their output format

#### 14.14.1.297. `.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].fieldRef`

Description

ObjectFieldSelector selects an APIVersioned field of an object.

Type

**object**

Required

- **fieldPath**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	Version of the schema the FieldPath is written in terms of, defaults to "v1".
<b>fieldPath</b>	<b>string</b>	Path of the field to select in the specified API version.

#### 14.14.1.298. `.spec.template.spec.volumes[].projected.sources[].downwardAPI.items[].resource`

Description

ResourceFieldSelector represents container resources (cpu, memory) and their output format

Type

**object**

Required

- **resource**

Property	Type	Description
<b>containerName</b>	<b>string</b>	Container name: required for volumes, optional for env vars
<b>divisor</b>	<b>Quantity</b>	Specifies the output format of the exposed resources, defaults to "1"
<b>resource</b>	<b>string</b>	Required: resource to select

#### 14.14.1.299. `.spec.template.spec.volumes[].projected.sources[].secret`

Description

Adapts a secret into a projected volume.

The contents of the target Secret's Data field will be presented in a projected volume as files using the keys in the Data field as the file names. Note that this is identical to a secret volume source without the default mode.

Type

**object**

Property	Type	Description
<b>items</b>	<b>array</b>	items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its key must be defined

#### 14.14.1.300. `.spec.template.spec.volumes[].projected.sources[].secret.items`

##### Description

items if unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

##### Type

**array**

#### 14.14.1.301. `.spec.template.spec.volumes[].projected.sources[].secret.items[]`

##### Description

Maps a string key to a path within a volume.

##### Type

**object**

##### Required

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

#### 14.14.1.302. `.spec.template.spec.volumes[].projected.sources[].serviceAccountToken`

##### Description

ServiceAccountTokenProjection represents a projected service account token volume. This projection can be used to insert a service account token into the pods runtime filesystem for use against APIs (Kubernetes API Server or otherwise).

##### Type

**object**

##### Required

- **path**

Property	Type	Description
<b>audience</b>	<b>string</b>	audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.

Property	Type	Description
<b>expirationSeconds</b>	<b>integer</b>	expirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
<b>path</b>	<b>string</b>	path is the path relative to the mount point of the file to project the token into.

#### 14.14.1.303. .spec.template.spec.volumes[].quobyte

##### Description

Represents a Quobyte mount that lasts the lifetime of a pod. Quobyte volumes do not support ownership management or SELinux relabeling.

##### Type

**object**

##### Required

- **registry**
- **volume**

Property	Type	Description
<b>group</b>	<b>string</b>	group to map volume access to Default is no group
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.



Property	Type	Description
<b>registry</b>	<b>string</b>	registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
<b>tenant</b>	<b>string</b>	tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
<b>user</b>	<b>string</b>	user to map volume access to Defaults to serviceaccount user
<b>volume</b>	<b>string</b>	volume is a string that references an already created Quobyte volume by name.

#### 14.14.1.304. .spec.template.spec.volumes[].rbd

##### Description

Represents a Rados Block Device mount that lasts the lifetime of a pod. RBD volumes support ownership management and SELinux relabeling.

##### Type

**object**

##### Required

- **monitors**
- **image**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#rbd">https://kubernetes.io/docs/concepts/storage/volumes#rbd</a>

Property	Type	Description
<b>image</b>	<b>string</b>	image is the rados image name. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>keyring</b>	<b>string</b>	keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>monitors</b>	<b>array (string)</b>	monitors is a collection of Ceph monitors. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>pool</b>	<b>string</b>	pool is the rados pool name. Default is rbd. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>readOnly</b>	<b>boolean</b>	readOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>user</b>	<b>string</b>	user is the rados user name. Default is admin. More info: <a href="https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it">https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it</a>

#### 14.14.1.305. .spec.template.spec.volumes[].rbd.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.14.1.306. .spec.template.spec.volumes[].scaleIO

##### Description

ScaleIOVolumeSource represents a persistent ScaleIO volume

##### Type

**object**

##### Required

- **gateway**
- **system**
- **secretRef**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
<b>gateway</b>	<b>string</b>	gateway is the host address of the ScaleIO API Gateway.
<b>protectionDomain</b>	<b>string</b>	protectionDomain is the name of the ScaleIO Protection Domain for the configured storage.
<b>readOnly</b>	<b>boolean</b>	readOnly Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>sslEnabled</b>	<b>boolean</b>	sslEnabled Flag enable/disable SSL communication with Gateway, default false

Property	Type	Description
<b>storageMode</b>	<b>string</b>	storageMode indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
<b>storagePool</b>	<b>string</b>	storagePool is the ScaleIO Storage Pool associated with the protection domain.
<b>system</b>	<b>string</b>	system is the name of the storage system as configured in ScaleIO.
<b>volumeName</b>	<b>string</b>	volumeName is the name of a volume already created in the ScaleIO system that is associated with this volume source.

#### 14.14.1.307. .spec.template.spec.volumes[].scaleIO.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.14.1.308. .spec.template.spec.volumes[].secret

##### Description

Adapts a Secret into a volume.

The contents of the target Secret's Data field will be presented in a volume as files using the keys in the Data field as the file names. Secret volumes support ownership management and SELinux relabeling.

##### Type

**object**

Property	Type	Description
<b>defaultMode</b>	<b>integer</b>	defaultMode is Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>items</b>	<b>array</b>	items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
<b>items[]</b>	<b>object</b>	Maps a string key to a path within a volume.
<b>optional</b>	<b>boolean</b>	optional field specify whether the Secret or its keys must be defined
<b>secretName</b>	<b>string</b>	secretName is the name of the secret in the pod's namespace to use. More info: <a href="https://kubernetes.io/docs/concepts/storage/volumes#secret">https://kubernetes.io/docs/concepts/storage/volumes#secret</a>

#### 14.14.1.309. .spec.template.spec.volumes[].secret.items

##### Description

items If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys

will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

**Type****array****14.14.1.310. .spec.template.spec.volumes[].secret.items[]****Description**

Maps a string key to a path within a volume.

**Type****object****Required**

- **key**
- **path**

Property	Type	Description
<b>key</b>	<b>string</b>	key is the key to project.
<b>mode</b>	<b>integer</b>	mode is Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
<b>path</b>	<b>string</b>	path is the relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

**14.14.1.311. .spec.template.spec.volumes[].storageos****Description**

Represents a StorageOS persistent volume resource.

**Type****object**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>readOnly</b>	<b>boolean</b>	readOnly defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
<b>secretRef</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>volumeName</b>	<b>string</b>	volumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
<b>volumeNamespace</b>	<b>string</b>	volumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

#### 14.14.1.312. .spec.template.spec.volumes[].storageos.secretRef

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

#### 14.14.1.313. .spec.template.spec.volumes[].vsphereVolume

##### Description

Represents a vSphere volume resource.

##### Type

**object**

##### Required

- **volumePath**

Property	Type	Description
<b>fsType</b>	<b>string</b>	fsType is filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
<b>storagePolicyID</b>	<b>string</b>	storagePolicyID is the storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
<b>storagePolicyName</b>	<b>string</b>	storagePolicyName is the storage Policy Based Management (SPBM) profile name.
<b>volumePath</b>	<b>string</b>	volumePath is the path that identifies vSphere volume vmdk

#### 14.14.1.314. .status

##### Description

ReplicationControllerStatus represents the current status of a replication controller.

##### Type

**object**

##### Required



- replicas

Property	Type	Description
<b>availableReplicas</b>	<b>integer</b>	The number of available replicas (ready for at least <code>minReadySeconds</code> ) for this replication controller.
<b>conditions</b>	<b>array</b>	Represents the latest available observations of a replication controller's current state.
<b>conditions[]</b>	<b>object</b>	<code>ReplicationControllerCondition</code> describes the state of a replication controller at a certain point.
<b>fullyLabeledReplicas</b>	<b>integer</b>	The number of pods that have labels matching the labels of the pod template of the replication controller.
<b>observedGeneration</b>	<b>integer</b>	<code>ObservedGeneration</code> reflects the generation of the most recently observed replication controller.
<b>readyReplicas</b>	<b>integer</b>	The number of ready replicas for this replication controller.
<b>replicas</b>	<b>integer</b>	Replicas is the most recently observed number of replicas. More info: <a href="https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#what-is-a-replicationcontroller">https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#what-is-a-replicationcontroller</a>

#### 14.14.1.315. `.status.conditions`

##### Description

Represents the latest available observations of a replication controller's current state.

##### Type

**array**

#### 14.14.1.316. `.status.conditions[]`

##### Description

`ReplicationControllerCondition` describes the state of a replication controller at a certain point.

## Type

object

## Required

- **type**
- **status**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	The last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	A human readable message indicating details about the transition.
<b>reason</b>	<b>string</b>	The reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of replication controller condition.

### 14.14.2. API endpoints

The following API endpoints are available:

- **/api/v1/replicationcontrollers**
  - **GET**: list or watch objects of kind ReplicationController
- **/api/v1/watch/replicationcontrollers**
  - **GET**: watch individual changes to a list of ReplicationController. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/replicationcontrollers**
  - **DELETE**: delete collection of ReplicationController
  - **GET**: list or watch objects of kind ReplicationController
  - **POST**: create a ReplicationController
- **/api/v1/watch/namespaces/{namespace}/replicationcontrollers**
  - **GET**: watch individual changes to a list of ReplicationController. deprecated: use the 'watch' parameter with a list operation instead.

- **/api/v1/namespaces/{namespace}/replicationcontrollers/{name}**
  - **DELETE:** delete a ReplicationController
  - **GET:** read the specified ReplicationController
  - **PATCH:** partially update the specified ReplicationController
  - **PUT:** replace the specified ReplicationController
- **/api/v1/watch/namespaces/{namespace}/replicationcontrollers/{name}**
  - **GET:** watch changes to an object of kind ReplicationController. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/namespaces/{namespace}/replicationcontrollers/{name}/status**
  - **GET:** read status of the specified ReplicationController
  - **PATCH:** partially update status of the specified ReplicationController
  - **PUT:** replace status of the specified ReplicationController

#### 14.14.2.1. /api/v1/replicationcontrollers

Table 14.374. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind ReplicationController

**Table 14.375. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">ReplicationControllerList</a> schema

HTTP code	Response body
401 - Unauthorized	Empty

#### 14.14.2.2. /api/v1/watch/replicationcontrollers

Table 14.376. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ReplicationController. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.377. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.14.2.3. /api/v1/namespaces/{namespace}/replicationcontrollers

Table 14.378. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.379. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ReplicationController

Table 14.380. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.381. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.382. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ReplicationController

Table 14.383. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.384. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ReplicationControllerList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a ReplicationController

**Table 14.385. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.386. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ReplicationController</b> schema	

**Table 14.387. HTTP responses**

HTTP code	Response body
200 - OK	<b>ReplicationController</b> schema

HTTP code	Response body
201 - Created	<a href="#">ReplicationController</a> schema
202 - Accepted	<a href="#">ReplicationController</a> schema
401 - Unauthorized	Empty

#### 14.14.2.4. /api/v1/watch/namespaces/{namespace}/replicationcontrollers

Table 14.388. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.389. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ReplicationController. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.390. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.14.2.5. /api/v1/namespaces/{namespace}/replicationcontrollers/{name}

Table 14.391. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ReplicationController
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.392. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a ReplicationController

Table 14.393. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.394. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.395. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ReplicationController

Table 14.396. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicationController</b> schema



HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ReplicationController

**Table 14.397. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.398. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.399. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicationController</b> schema
201 - Created	<b>ReplicationController</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ReplicationController

Table 14.400. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.401. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ReplicationController</b> schema	

Table 14.402. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicationController</b> schema
201 - Created	<b>ReplicationController</b> schema
401 - Unauthorized	Empty

#### 14.14.2.6. /api/v1/watch/namespaces/{namespace}/replicationcontrollers/{name}

Table 14.403. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ReplicationController
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 14.404. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ReplicationController. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.405. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.14.2.7. /api/v1/namespaces/{namespace}/replicationcontrollers/{name}/status

Table 14.406. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ReplicationController
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.407. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified ReplicationController

Table 14.408. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">ReplicationController</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified ReplicationController

Table 14.409. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.410. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.411. HTTP responses



HTTP code	Response body
200 - OK	<a href="#">ReplicationController</a> schema
201 - Created	<a href="#">ReplicationController</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified ReplicationController

**Table 14.412. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.413. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ReplicationController</b> r schema	

Table 14.414. HTTP responses

HTTP code	Response body
200 - OK	<b>ReplicationController</b> schema
201 - Created	<b>ReplicationController</b> schema
401 - Unauthorized	Empty

## 14.15. RESOURCEQUOTA [V1]

### Description

ResourceQuota sets aggregate quota restrictions enforced per namespace

### Type

**object**

### 14.15.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	ResourceQuotaSpec defines the desired hard limits to enforce for Quota.
<b>status</b>	<b>object</b>	ResourceQuotaStatus defines the enforced hard limits and observed use.

#### 14.15.1.1. .spec

##### Description

ResourceQuotaSpec defines the desired hard limits to enforce for Quota.

##### Type

**object**

Property	Type	Description
<b>hard</b>	<b>object (Quantity)</b>	hard is the set of desired hard limits for each named resource. More info: <a href="https://kubernetes.io/docs/concepts/policy/resource-quotas/">https://kubernetes.io/docs/concepts/policy/resource-quotas/</a>
<b>scopeSelector</b>	<b>object</b>	A scope selector represents the AND of the selectors represented by the scoped-resource selector requirements.
<b>scopes</b>	<b>array (string)</b>	A collection of filters that must match each object tracked by a quota. If not specified, the quota matches all objects.

#### 14.15.1.2. .spec.scopeSelector

##### Description

A scope selector represents the AND of the selectors represented by the scoped-resource selector requirements.

##### Type

**object**

Property	Type	Description
<b>matchExpressions</b>	<b>array</b>	A list of scope selector requirements by scope of the resources.
<b>matchExpressions[]</b>	<b>object</b>	A scoped-resource selector requirement is a selector that contains values, a scope name, and an operator that relates the scope name and values.

**14.15.1.3. .spec.scopeSelector.matchExpressions****Description**

A list of scope selector requirements by scope of the resources.

**Type**

**array**

**14.15.1.4. .spec.scopeSelector.matchExpressions[]****Description**

A scoped-resource selector requirement is a selector that contains values, a scope name, and an operator that relates the scope name and values.

**Type**

**object**

**Required**

- **scopeName**
- **operator**

Property	Type	Description
<b>operator</b>	<b>string</b>	Represents a scope's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist.  Possible enum values: - <b>"DoesNotExist"</b> - <b>"Exists"</b> - <b>"In"</b> - <b>"NotIn"</b>

Property	Type	Description
<b>scopeName</b>	<b>string</b>	<p>The name of the scope that the selector applies to.</p> <p>Possible enum values: -  <b>"BestEffort"</b> Match all pod objects that have best effort quality of service -  <b>"CrossNamespacePodAffinity"</b> Match all pod objects that have cross-namespace pod (anti)affinity mentioned. -  <b>"NotBestEffort"</b> Match all pod objects that do not have best effort quality of service -  <b>"NotTerminating"</b> Match all pod objects where spec.activeDeadlineSeconds is nil -  <b>"PriorityClass"</b> Match all pod objects that have priority class mentioned -  <b>"Terminating"</b> Match all pod objects where spec.activeDeadlineSeconds &gt;=0</p>
<b>values</b>	<b>array (string)</b>	<p>An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.</p>

#### 14.15.1.5. .status

##### Description

ResourceQuotaStatus defines the enforced hard limits and observed use.

##### Type

**object**

Property	Type	Description
<b>hard</b>	<b>object (Quantity)</b>	<p>Hard is the set of enforced hard limits for each named resource. More info: <a href="https://kubernetes.io/docs/concepts/policy/resource-quotas/">https://kubernetes.io/docs/concepts/policy/resource-quotas/</a></p>

Property	Type	Description
<b>used</b>	<b>object (Quantity)</b>	Used is the current observed total usage of the resource in the namespace.

## 14.15.2. API endpoints

The following API endpoints are available:

- **/api/v1/resourcequotas**
  - **GET**: list or watch objects of kind ResourceQuota
- **/api/v1/watch/resourcequotas**
  - **GET**: watch individual changes to a list of ResourceQuota. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/resourcequotas**
  - **DELETE**: delete collection of ResourceQuota
  - **GET**: list or watch objects of kind ResourceQuota
  - **POST**: create a ResourceQuota
- **/api/v1/watch/namespaces/{namespace}/resourcequotas**
  - **GET**: watch individual changes to a list of ResourceQuota. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/resourcequotas/{name}**
  - **DELETE**: delete a ResourceQuota
  - **GET**: read the specified ResourceQuota
  - **PATCH**: partially update the specified ResourceQuota
  - **PUT**: replace the specified ResourceQuota
- **/api/v1/watch/namespaces/{namespace}/resourcequotas/{name}**
  - **GET**: watch changes to an object of kind ResourceQuota. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/namespaces/{namespace}/resourcequotas/{name}/status**
  - **GET**: read status of the specified ResourceQuota
  - **PATCH**: partially update status of the specified ResourceQuota
  - **PUT**: replace status of the specified ResourceQuota

## 14.15.2.1. /api/v1/resourcequotas

Table 14.415. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind ResourceQuota

Table 14.416. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ResourceQuotaList</a> schema
401 - Unauthorized	Empty

### 14.15.2.2. /api/v1/watch/resourcequotas

Table 14.417. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ResourceQuota. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.418. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.15.2.3. /api/v1/namespaces/{namespace}/resourcequotas

Table 14.419. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.420. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ResourceQuota

Table 14.421. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>"k8s.io/initial-events-end": "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>



Table 14.422. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.423. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ResourceQuota

Table 14.424. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.425. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ResourceQuotaList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a ResourceQuota

**Table 14.426. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.427. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ResourceQuota</b> schema	

**Table 14.428. HTTP responses**

HTTP code	Response body
200 - OK	<b>ResourceQuota</b> schema

HTTP code	Response body
201 - Created	<a href="#">ResourceQuota</a> schema
202 - Accepted	<a href="#">ResourceQuota</a> schema
401 - Unauthorized	Empty

#### 14.15.2.4. /api/v1/watch/namespaces/{namespace}/resourcequotas

Table 14.429. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.430. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ResourceQuota. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.431. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.15.2.5. /api/v1/namespaces/{namespace}/resourcequotas/{name}

Table 14.432. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ResourceQuota
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.433. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a ResourceQuota

Table 14.434. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.435. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.436. HTTP responses

HTTP code	Response body
200 - OK	<b>ResourceQuota</b> schema
202 - Accepted	<b>ResourceQuota</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ResourceQuota

Table 14.437. HTTP responses

HTTP code	Response body
200 - OK	<b>ResourceQuota</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ResourceQuota

**Table 14.438. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.439. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.440. HTTP responses

HTTP code	Reponse body
200 - OK	<b>ResourceQuota</b> schema
201 - Created	<b>ResourceQuota</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ResourceQuota

Table 14.441. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.442. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ResourceQuota</b> schema	

Table 14.443. HTTP responses

HTTP code	Response body
200 - OK	<b>ResourceQuota</b> schema
201 - Created	<b>ResourceQuota</b> schema
401 - Unauthorized	Empty

#### 14.15.2.6. /api/v1/watch/namespaces/{namespace}/resourcequotas/{name}

Table 14.444. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ResourceQuota
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 14.445. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ResourceQuota. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.446. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.15.2.7. /api/v1/namespaces/{namespace}/resourcequotas/{name}/status

Table 14.447. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ResourceQuota
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.448. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified ResourceQuota

Table 14.449. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ResourceQuota</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified ResourceQuota

Table 14.450. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.451. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.452. HTTP responses

HTTP code	Response body
200 - OK	<b>ResourceQuota</b> schema
201 - Created	<b>ResourceQuota</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified ResourceQuota

**Table 14.453. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.454. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ResourceQuota</b> schema	

Table 14.455. HTTP responses

HTTP code	Response body
200 - OK	<b>ResourceQuota</b> schema
201 - Created	<b>ResourceQuota</b> schema
401 - Unauthorized	Empty

## 14.16. SECRET [V1]

### Description

Secret holds secret data of a certain type. The total bytes of the values in the Data field must be less than MaxSecretSize bytes.

### Type

**object**

### 14.16.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>data</b>	<b>object (string)</b>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in <a href="https://tools.ietf.org/html/rfc4648#section-4">https://tools.ietf.org/html/rfc4648#section-4</a>

Property	Type	Description
<b>immutable</b>	<b>boolean</b>	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>stringData</b>	<b>object (string)</b>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only input field for convenience. All keys and values are merged into the data field on write, overwriting any existing values. The stringData field is never output when reading from the API.
<b>type</b>	<b>string</b>	Used to facilitate programmatic handling of secret data. More info: <a href="https://kubernetes.io/docs/concepts/configuration/secret/#secret-types">https://kubernetes.io/docs/concepts/configuration/secret/#secret-types</a>

### 14.16.2. API endpoints

The following API endpoints are available:

- **/api/v1/secrets**
  - **GET**: list or watch objects of kind Secret

- **/api/v1/watch/secrets**
  - **GET**: watch individual changes to a list of Secret. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/secrets**
  - **DELETE**: delete collection of Secret
  - **GET**: list or watch objects of kind Secret
  - **POST**: create a Secret
- **/api/v1/watch/namespaces/{namespace}/secrets**
  - **GET**: watch individual changes to a list of Secret. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/secrets/{name}**
  - **DELETE**: delete a Secret
  - **GET**: read the specified Secret
  - **PATCH**: partially update the specified Secret
  - **PUT**: replace the specified Secret
- **/api/v1/watch/namespaces/{namespace}/secrets/{name}**
  - **GET**: watch changes to an object of kind Secret. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 14.16.2.1. /api/v1/secrets

Table 14.456. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Secret

**Table 14.457. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">SecretList</a> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 14.16.2.2. /api/v1/watch/secrets

Table 14.458. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Secret. deprecated: use the 'watch' parameter with a list operation instead.

**Table 14.459. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.16.2.3. /api/v1/namespaces/{namespace}/secrets

Table 14.460. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.461. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Secret

Table 14.462. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.463. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.464. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Secret

Table 14.465. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.466. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">SecretList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a Secret

**Table 14.467. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.468. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Secret</b> schema	

**Table 14.469. HTTP responses**

HTTP code	Response body
200 - OK	<b>Secret</b> schema

HTTP code	Response body
201 - Created	<b>Secret</b> schema
202 - Accepted	<b>Secret</b> schema
401 - Unauthorized	Empty

#### 14.16.2.4. /api/v1/watch/namespaces/{namespace}/secrets

Table 14.470. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.471. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Secret. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.472. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.16.2.5. /api/v1/namespaces/{namespace}/secrets/{name}

Table 14.473. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Secret
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.474. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a Secret

Table 14.475. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.476. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.477. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Secret

Table 14.478. HTTP responses

HTTP code	Response body
200 - OK	<b>Secret</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Secret

**Table 14.479. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.480. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.481. HTTP responses

HTTP code	Response body
200 - OK	<b>Secret</b> schema
201 - Created	<b>Secret</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Secret

Table 14.482. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.483. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Secret</b> schema	

Table 14.484. HTTP responses

HTTP code	Response body
200 - OK	<b>Secret</b> schema
201 - Created	<b>Secret</b> schema
401 - Unauthorized	Empty

#### 14.16.2.6. /api/v1/watch/namespaces/{namespace}/secrets/{name}

Table 14.485. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Secret
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.486. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Secret. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.487. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 14.17. SERVICE [V1]

### Description

Service is a named abstraction of software service (for example, mysql) consisting of local port (for example 3306) that the proxy listens on, and the selector that determines which pods will answer requests sent through the proxy.

### Type

**object**

### 14.17.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>spec</b>	<b>object</b>	ServiceSpec describes the attributes that a user creates on a service.
<b>status</b>	<b>object</b>	ServiceStatus represents the current status of a service.

#### 14.17.1.1. .spec

##### Description

ServiceSpec describes the attributes that a user creates on a service.

##### Type

**object**

Property	Type	Description
<b>allocateLoadBalancerNodePorts</b>	<b>boolean</b>	allocateLoadBalancerNodePorts defines if NodePorts will be automatically allocated for services with type LoadBalancer. Default is "true". It may be set to "false" if the cluster load-balancer does not rely on NodePorts. If the caller requests specific NodePorts (by specifying a value), those requests will be respected, regardless of this field. This field may only be set for services with type LoadBalancer and will be cleared if the type is changed to any other type.

Property	Type	Description
<b>clusterIP</b>	<b>string</b>	<p>clusterIP is the IP address of the service and is usually assigned randomly. If an address is specified manually, is in-range (as per system configuration), and is not in use, it will be allocated to the service; otherwise creation of the service will fail. This field may not be changed through updates unless the type field is also being changed to ExternalName (which requires this field to be blank) or the type field is being changed from ExternalName (in which case this field may optionally be specified, as describe above). Valid values are "None", empty string (""), or a valid IP address. Setting this to "None" makes a "headless service" (no virtual IP), which is useful when direct endpoint connections are preferred and proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. If this field is specified when creating a Service of type ExternalName, creation will fail. This field will be wiped when updating a Service to type ExternalName. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies">https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies</a></p>

Property	Type	Description
<b>clusterIPs</b>	<b>array (string)</b>	<p>ClusterIPs is a list of IP addresses assigned to this service, and are usually assigned randomly. If an address is specified manually, is in-range (as per system configuration), and is not in use, it will be allocated to the service; otherwise creation of the service will fail. This field may not be changed through updates unless the type field is also being changed to ExternalName (which requires this field to be empty) or the type field is being changed from ExternalName (in which case this field may optionally be specified, as describe above). Valid values are "None", empty string (""), or a valid IP address. Setting this to "None" makes a "headless service" (no virtual IP), which is useful when direct endpoint connections are preferred and proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. If this field is specified when creating a Service of type ExternalName, creation will fail. This field will be wiped when updating a Service to type ExternalName. If this field is not specified, it will be initialized from the clusterIP field. If this field is specified, clients must ensure that clusterIPs[0] and clusterIP have the same value.</p> <p>This field may hold a maximum of two entries (dual-stack IPs, in either order). These IPs must correspond to the values of the ipFamilies field. Both clusterIPs and ipFamilies are governed by the ipFamilyPolicy field. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies">https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies</a></p>

Property	Type	Description
<b>externalIPs</b>	<b>array (string)</b>	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
<b>externalName</b>	<b>string</b>	externalName is the external reference that discovery mechanisms will return as an alias for this service (e.g. a DNS CNAME record). No proxying will be involved. Must be a lowercase RFC-1123 hostname ( <a href="https://tools.ietf.org/html/rfc1123">https://tools.ietf.org/html/rfc1123</a> ) and requires <b>type</b> to be "ExternalName".

Property	Type	Description
<b>externalTrafficPolicy</b>	<b>string</b>	<p>externalTrafficPolicy describes how nodes distribute service traffic they receive on one of the Service's "externally-facing" addresses (NodePorts, ExternalIPs, and LoadBalancer IPs). If set to "Local", the proxy will configure the service in a way that assumes that external load balancers will take care of balancing the service traffic between nodes, and so each node will deliver traffic only to the node-local endpoints of the service, without masquerading the client source IP. (Traffic mistakenly sent to a node with no endpoints will be dropped.) The default value, "Cluster", uses the standard behavior of routing to all endpoints evenly (possibly modified by topology and other features). Note that traffic sent to an External IP or LoadBalancer IP from within the cluster will always get "Cluster" semantics, but clients sending to a NodePort from within the cluster may need to take traffic policy into account when picking a node.</p> <p>Possible enum values: - <b>"Cluster"</b> - <b>"Cluster"</b> routes traffic to all endpoints. - <b>"Local"</b> - <b>"Local"</b> preserves the source IP of the traffic by routing only to endpoints on the same node as the traffic was received on (dropping the traffic if there are no local endpoints).</p>



Property	Type	Description
<b>healthCheckNodePort</b>	<b>integer</b>	<p>healthCheckNodePort specifies the healthcheck nodePort for the service. This only applies when type is set to LoadBalancer and externalTrafficPolicy is set to Local. If a value is specified, is in-range, and is not in use, it will be used. If not specified, a value will be automatically allocated. External systems (e.g. load-balancers) can use this port to determine if a given node holds endpoints for this service or not. If this field is specified when creating a Service which does not need it, creation will fail. This field will be wiped when updating a Service to no longer need it (e.g. changing type). This field cannot be updated once set.</p>
<b>internalTrafficPolicy</b>	<b>string</b>	<p>InternalTrafficPolicy describes how nodes distribute service traffic they receive on the ClusterIP. If set to "Local", the proxy will assume that pods only want to talk to endpoints of the service on the same node as the pod, dropping the traffic if there are no local endpoints. The default value, "Cluster", uses the standard behavior of routing to all endpoints evenly (possibly modified by topology and other features).</p> <p>Possible enum values: - <b>"Cluster"</b> routes traffic to all endpoints. - <b>"Local"</b> routes traffic only to endpoints on the same node as the client pod (dropping the traffic if there are no local endpoints).</p>

Property	Type	Description
<b>ipFamilies</b>	<b>array (string)</b>	<p>IPFamilies is a list of IP families (e.g. IPv4, IPv6) assigned to this service. This field is usually assigned automatically based on cluster configuration and the ipFamilyPolicy field. If this field is specified manually, the requested family is available in the cluster, and ipFamilyPolicy allows it, it will be used; otherwise creation of the service will fail. This field is conditionally mutable: it allows for adding or removing a secondary IP family, but it does not allow changing the primary IP family of the Service. Valid values are "IPv4" and "IPv6". This field only applies to Services of types ClusterIP, NodePort, and LoadBalancer, and does apply to "headless" services. This field will be wiped when updating a Service to type ExternalName.</p> <p>This field may hold a maximum of two entries (dual-stack families, in either order). These families must correspond to the values of the clusterIPs field, if specified. Both clusterIPs and ipFamilies are governed by the ipFamilyPolicy field.</p>
<b>ipFamilyPolicy</b>	<b>string</b>	<p>IPFamilyPolicy represents the dual-stack-ness requested or required by this Service. If there is no value provided, then this field will be set to SingleStack. Services can be "SingleStack" (a single IP family), "PreferDualStack" (two IP families on dual-stack configured clusters or a single IP family on single-stack clusters), or "RequireDualStack" (two IP families on dual-stack configured clusters, otherwise fail). The ipFamilies and clusterIPs fields depend on the value of this field. This field will be wiped when updating a service to type ExternalName.</p>

Property	Type	Possible enum values: - Description
		<p>that this service prefers dual-stack when the cluster is configured for dual-stack. If the cluster is not configured for dual-stack the service will be assigned a single IPFamily. If the IPFamily is not set in service.spec.ipFamilies then the service will be assigned the default IPFamily configured on the cluster -</p> <p><b>"RequireDualStack"</b> indicates that this service requires dual-stack. Using IPFamilyPolicyRequireDualStack on a single stack cluster will result in validation errors. The IPFamilies (and their order) assigned to this service is based on service.spec.ipFamilies. If service.spec.ipFamilies was not provided then it will be assigned according to how they are configured on the cluster. If service.spec.ipFamilies has only one entry then the alternative IPFamily will be added by apiserver - <b>"SingleStack"</b> indicates that this service is required to have a single IPFamily. The IPFamily assigned is based on the default IPFamily used by the cluster or as identified by service.spec.ipFamilies field</p>

Property	Type	Description
<b>loadBalancerClass</b>	<b>string</b>	<p>loadBalancerClass is the class of the load balancer implementation this Service belongs to. If specified, the value of this field must be a label-style identifier, with an optional prefix, e.g. "internal-vip" or "example.com/internal-vip". Unprefixed names are reserved for end-users. This field can only be set when the Service type is 'LoadBalancer'. If not set, the default load balancer implementation is used, today this is typically done through the cloud provider integration, but should apply for any default implementation. If set, it is assumed that a load balancer implementation is watching for Services with a matching class. Any default load balancer implementation (e.g. cloud providers) should ignore Services that set this field. This field can only be set when creating or updating a Service to type 'LoadBalancer'. Once set, it can not be changed. This field will be wiped when a service is updated to a non 'LoadBalancer' type.</p>
<b>loadBalancerIP</b>	<b>string</b>	<p>Only applies to Service Type: LoadBalancer. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature. Deprecated: This field was under-specified and its meaning varies across implementations, and it cannot support dual-stack. As of Kubernetes v1.24, users are encouraged to use implementation-specific annotations when available. This field may be removed in a future API version.</p>

Property	Type	Description
<b>loadBalancerSourceRanges</b>	<b>array (string)</b>	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: <a href="https://kubernetes.io/docs/tasks/access-application-cluster/create-external-load-balancer/">https://kubernetes.io/docs/tasks/access-application-cluster/create-external-load-balancer/</a>
<b>ports</b>	<b>array</b>	The list of ports that are exposed by this service. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies">https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies</a>
<b>ports[]</b>	<b>object</b>	ServicePort contains information on service's port.
<b>publishNotReadyAddresses</b>	<b>boolean</b>	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.

Property	Type	Description
<b>selector</b>	<b>object (string)</b>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>
<b>sessionAffinity</b>	<b>string</b>	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies">https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies</a>  Possible enum values: - <b>"ClientIP"</b> is the Client IP based. - <b>"None"</b> - no session affinity.
<b>sessionAffinityConfig</b>	<b>object</b>	SessionAffinityConfig represents the configurations of session affinity.

Property	Type	Description
<b>type</b>	<b>string</b>	<p>type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object or EndpointSlice objects. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a virtual IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the same endpoints as the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the same endpoints as the clusterIP. "ExternalName" aliases this service to the specified externalName. Several other fields do not apply to ExternalName services. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types">https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types</a></p> <p>Possible enum values: -  <b>"ClusterIP"</b> means a service will only be accessible inside the cluster, via the cluster IP. -  <b>"ExternalName"</b> means a service consists of only a reference to an external name that kubedns or equivalent will return as a CNAME record, with no exposing or proxying of any pods involved. -  <b>"LoadBalancer"</b> means a service will be exposed via an external load balancer (if the cloud provider supports it), in addition to 'NodePort' type. -  <b>"NodePort"</b> means a service will</p>

Property	Type	Description
		be exposed on one port of every node, in addition to 'ClusterIP' type.

### 14.17.1.2. .spec.ports

#### Description

The list of ports that are exposed by this service. More info:

<https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies>

#### Type

**array**

### 14.17.1.3. .spec.ports[]

#### Description

ServicePort contains information on service's port.

#### Type

**object**

#### Required

- **port**

Property	Type	Description
<b>appProtocol</b>	<b>string</b>	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and <a href="https://www.iana.org/assignments/service-names">https://www.iana.org/assignments/service-names</a> ). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol.
<b>name</b>	<b>string</b>	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.



Property	Type	Description
<b>nodePort</b>	<b>integer</b>	<p>The port on each node on which this service is exposed when type is NodePort or LoadBalancer. Usually assigned by the system. If a value is specified, in-range, and not in use it will be used, otherwise the operation will fail. If not specified, a port will be allocated if this Service requires one. If this field is specified when creating a Service which does not need it, creation will fail. This field will be wiped when updating a Service to no longer need it (e.g. changing type from NodePort to ClusterIP). More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport">https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport</a></p>
<b>port</b>	<b>integer</b>	The port that will be exposed by this service.
<b>protocol</b>	<b>string</b>	<p>The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.</p> <p>Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.</p>
<b>targetPort</b>	<b>IntOrString</b>	<p>Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: <a href="https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service">https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service</a></p>

#### 14.17.1.4. .spec.sessionAffinityConfig

##### Description

SessionAffinityConfig represents the configurations of session affinity.

##### Type

**object**

Property	Type	Description
<b>clientIP</b>	<b>object</b>	ClientIPConfig represents the configurations of Client IP based session affinity.

#### 14.17.1.5. .spec.sessionAffinityConfig.clientIP

##### Description

ClientIPConfig represents the configurations of Client IP based session affinity.

##### Type

**object**

Property	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && ≤86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

#### 14.17.1.6. .status

##### Description

ServiceStatus represents the current status of a service.

##### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array (Condition)</b>	Current service state
<b>loadBalancer</b>	<b>object</b>	LoadBalancerStatus represents the status of a load-balancer.

#### 14.17.1.7. .status.loadBalancer

##### Description

LoadBalancerStatus represents the status of a load-balancer.

#### Type

**object**

Property	Type	Description
<b>ingress</b>	<b>array</b>	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.
<b>ingress[]</b>	<b>object</b>	LoadBalancerIngress represents the status of a load-balancer ingress point: traffic intended for the service should be sent to an ingress point.

#### 14.17.1.8. .status.loadBalancer.ingress

##### Description

Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

##### Type

**array**

#### 14.17.1.9. .status.loadBalancer.ingress[]

##### Description

LoadBalancerIngress represents the status of a load-balancer ingress point: traffic intended for the service should be sent to an ingress point.

##### Type

**object**

Property	Type	Description
<b>hostname</b>	<b>string</b>	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
<b>ip</b>	<b>string</b>	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

Property	Type	Description
<b>ports</b>	<b>array</b>	Ports is a list of records of service ports If used, every port defined in the service should have an entry in it
<b>ports[]</b>	<b>object</b>	

#### 14.17.1.10. .status.loadBalancer.ingress[].ports

##### Description

Ports is a list of records of service ports If used, every port defined in the service should have an entry in it

##### Type

**array**

#### 14.17.1.11. .status.loadBalancer.ingress[].ports[]

##### Description

##### Type

**object**

##### Required

- **port**
- **protocol**

Property	Type	Description
<b>error</b>	<b>string</b>	Error is to record the problem with the service port The format of the error shall comply with the following rules: - built-in error values shall be specified in this file and those shall use CamelCase names - cloud provider specific error values must have names that comply with the format foo.example.com/CamelCase.
<b>port</b>	<b>integer</b>	Port is the port number of the service port of which status is recorded here

Property	Type	Description
<b>protocol</b>	<b>string</b>	<p>Protocol is the protocol of the service port of which status is recorded here The supported values are: "TCP", "UDP", "SCTP"</p> <p>Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.</p>

### 14.17.2. API endpoints

The following API endpoints are available:

- **/api/v1/services**
  - **GET**: list or watch objects of kind Service
- **/api/v1/watch/services**
  - **GET**: watch individual changes to a list of Service. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/services**
  - **DELETE**: delete collection of Service
  - **GET**: list or watch objects of kind Service
  - **POST**: create a Service
- **/api/v1/watch/namespaces/{namespace}/services**
  - **GET**: watch individual changes to a list of Service. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/services/{name}**
  - **DELETE**: delete a Service
  - **GET**: read the specified Service
  - **PATCH**: partially update the specified Service
  - **PUT**: replace the specified Service
- **/api/v1/watch/namespaces/{namespace}/services/{name}**
  - **GET**: watch changes to an object of kind Service. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/api/v1/namespaces/{namespace}/services/{name}/status**
  - **GET**: read status of the specified Service

- **PATCH**: partially update status of the specified Service
- **PUT**: replace status of the specified Service

### 14.17.2.1. /api/v1/services

Table 14.488. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Service

Table 14.489. HTTP responses



HTTP code	Response body
200 - OK	<a href="#">ServiceList</a> schema
401 - Unauthorized	Empty

### 14.17.2.2. /api/v1/watch/services

Table 14.490. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Service. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.491. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.17.2.3. /api/v1/namespaces/{namespace}/services

Table 14.492. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.493. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Service

Table 14.494. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.495. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.496. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Service

Table 14.497. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.498. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ServiceList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a Service

**Table 14.499. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.500. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Service</b> schema	

**Table 14.501. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Service</b> schema

HTTP code	Reponse body
201 - Created	<b>Service</b> schema
202 - Accepted	<b>Service</b> schema
401 - Unauthorized	Empty

#### 14.17.2.4. /api/v1/watch/namespaces/{namespace}/services

Table 14.502. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.503. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Service. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.504. HTTP responses



HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.17.2.5. /api/v1/namespaces/{namespace}/services/{name}

Table 14.505. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Service
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.506. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a Service

Table 14.507. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.508. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.509. HTTP responses

HTTP code	Response body
200 - OK	<b>Service</b> schema
202 - Accepted	<b>Service</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Service

Table 14.510. HTTP responses

HTTP code	Response body
200 - OK	<b>Service</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Service

**Table 14.511. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.512. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.513. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Service</b> schema
201 - Created	<b>Service</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Service

Table 14.514. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.515. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Service</b> schema	

Table 14.516. HTTP responses

HTTP code	Response body
200 - OK	<b>Service</b> schema
201 - Created	<b>Service</b> schema
401 - Unauthorized	Empty

#### 14.17.2.6. /api/v1/watch/namespaces/{namespace}/services/{name}

Table 14.517. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Service
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.518. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Service. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 14.519. HTTP responses



HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.17.2.7. /api/v1/namespaces/{namespace}/services/{name}/status

Table 14.520. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Service
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.521. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Service

Table 14.522. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">Service</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Service

Table 14.523. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.524. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.525. HTTP responses

HTTP code	Response body
200 - OK	<b>Service</b> schema
201 - Created	<b>Service</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Service

**Table 14.526. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.527. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Service</b> schema	

Table 14.528. HTTP responses

HTTP code	Response body
200 - OK	<b>Service</b> schema
201 - Created	<b>Service</b> schema
401 - Unauthorized	Empty

## 14.18. SERVICEACCOUNT [V1]

### Description

ServiceAccount binds together: \* a name, understood by users, and perhaps by peripheral systems, for an identity \* a principal that can be authenticated and authorized \* a set of secrets

### Type

**object**

### 14.18.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>automountServiceAccountToken</b>	<b>boolean</b>	AutomountServiceAccountToken indicates whether pods running as this service account should have an API token automatically mounted. Can be overridden at the pod level.

Property	Type	Description
<b>imagePullSecrets</b>	<b>array</b>	ImagePullSecrets is a list of references to secrets in the same namespace to use for pulling any images in pods that reference this ServiceAccount. ImagePullSecrets are distinct from Secrets because Secrets can be mounted in the pod, but ImagePullSecrets are only accessed by the kubelet. More info: <a href="https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod">https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod</a>
<b>imagePullSecrets[]</b>	<b>object</b>	LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>secrets</b>	<b>array</b>	Secrets is a list of the secrets in the same namespace that pods running using this ServiceAccount are allowed to use. Pods are only limited to this list if this service account has a "kubernetes.io/enforce-mountable-secrets" annotation set to "true". This field should not be used to find auto-generated service account token secrets for use outside of pods. Instead, tokens can be requested directly using the TokenRequest API, or service account token secrets can be manually created. More info: <a href="https://kubernetes.io/docs/concepts/configuration/secret">https://kubernetes.io/docs/concepts/configuration/secret</a>
<b>secrets[]</b>	<b>object</b>	ObjectReference contains enough information to let you inspect or modify the referred object.

#### 14.18.1.1. .imagePullSecrets

##### Description

ImagePullSecrets is a list of references to secrets in the same namespace to use for pulling any images in pods that reference this ServiceAccount. ImagePullSecrets are distinct from Secrets because Secrets can be mounted in the pod, but ImagePullSecrets are only accessed by the kubelet. More info: <https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod>

##### Type

**array**

#### 14.18.1.2. .imagePullSecrets[]

##### Description

LocalObjectReference contains enough information to let you locate the referenced object inside the same namespace.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>

### 14.18.1.3. .secrets

#### Description

Secrets is a list of the secrets in the same namespace that pods running using this ServiceAccount are allowed to use. Pods are only limited to this list if this service account has a "kubernetes.io/enforce-mountable-secrets" annotation set to "true". This field should not be used to find auto-generated service account token secrets for use outside of pods. Instead, tokens can be requested directly using the TokenRequest API, or service account token secrets can be manually created. More info: <https://kubernetes.io/docs/concepts/configuration/secret>

#### Type

**array**

### 14.18.1.4. .secrets[]

#### Description

ObjectReference contains enough information to let you inspect or modify the referred object.

#### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>



## 14.18.2. API endpoints

The following API endpoints are available:

- **/api/v1/serviceaccounts**
  - **GET**: list or watch objects of kind ServiceAccount
- **/api/v1/watch/serviceaccounts**
  - **GET**: watch individual changes to a list of ServiceAccount. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/serviceaccounts**
  - **DELETE**: delete collection of ServiceAccount
  - **GET**: list or watch objects of kind ServiceAccount
  - **POST**: create a ServiceAccount
- **/api/v1/watch/namespaces/{namespace}/serviceaccounts**
  - **GET**: watch individual changes to a list of ServiceAccount. deprecated: use the 'watch' parameter with a list operation instead.
- **/api/v1/namespaces/{namespace}/serviceaccounts/{name}**
  - **DELETE**: delete a ServiceAccount
  - **GET**: read the specified ServiceAccount
  - **PATCH**: partially update the specified ServiceAccount
  - **PUT**: replace the specified ServiceAccount
- **/api/v1/watch/namespaces/{namespace}/serviceaccounts/{name}**
  - **GET**: watch changes to an object of kind ServiceAccount. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 14.18.2.1. /api/v1/serviceaccounts

Table 14.529. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind ServiceAccount

Table 14.530. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ServiceAccountList</a> schema
401 - Unauthorized	Empty

#### 14.18.2.2. /api/v1/watch/serviceaccounts

Table 14.531. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ServiceAccount. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.532. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 14.18.2.3. /api/v1/namespaces/{namespace}/serviceaccounts

Table 14.533. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.534. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ServiceAccount

Table 14.535. Query parameters

Parameter	Type	Description
-----------	------	-------------



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 14.536. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.537. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ServiceAccount

Table 14.538. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 14.539. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ServiceAccountList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a ServiceAccount

**Table 14.540. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 14.541. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ServiceAccount</b> schema	

**Table 14.542. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>ServiceAccount</b> schema



HTTP code	Reponse body
201 - Created	<b>ServiceAccount</b> schema
202 - Accepted	<b>ServiceAccount</b> schema
401 - Unauthorized	Empty

#### 14.18.2.4. /api/v1/watch/namespaces/{namespace}/serviceaccounts

Table 14.543. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.544. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ServiceAccount. deprecated: use the 'watch' parameter with a list operation instead.

Table 14.545. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 14.18.2.5. /api/v1/namespaces/{namespace}/serviceaccounts/{name}

Table 14.546. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ServiceAccount
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 14.547. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a ServiceAccount

Table 14.548. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 14.549. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 14.550. HTTP responses

HTTP code	Response body
200 - OK	<b>ServiceAccount</b> schema
202 - Accepted	<b>ServiceAccount</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ServiceAccount

Table 14.551. HTTP responses

HTTP code	Response body
200 - OK	<b>ServiceAccount</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ServiceAccount

**Table 14.552. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 14.553. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 14.554. HTTP responses

HTTP code	Reponse body
200 - OK	<b>ServiceAccount</b> schema
201 - Created	<b>ServiceAccount</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ServiceAccount

Table 14.555. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 14.556. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ServiceAccount</b> schema	

Table 14.557. HTTP responses

HTTP code	Response body
200 - OK	<b>ServiceAccount</b> schema
201 - Created	<b>ServiceAccount</b> schema
401 - Unauthorized	Empty

#### 14.18.2.6. /api/v1/watch/namespaces/{namespace}/serviceaccounts/{name}

Table 14.558. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ServiceAccount
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 14.559. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>"k8s.io/initial-events-end": "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ServiceAccount. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 14.560. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## CHAPTER 15. DISCOVERY APIS

### 15.1. DISCOVERY APIS

#### 15.1.1. EndpointSlice [discovery.k8s.io/v1]

##### Description

EndpointSlice represents a subset of the endpoints that implement a service. For a given service there may be multiple EndpointSlice objects, selected by labels, which must be joined to produce the full set of endpoints.

##### Type

**object**

### 15.2. ENDPPOINTSICE [DISCOVERY.K8S.IO/V1]

##### Description

EndpointSlice represents a subset of the endpoints that implement a service. For a given service there may be multiple EndpointSlice objects, selected by labels, which must be joined to produce the full set of endpoints.

##### Type

**object**

##### Required

- **addressType**
- **endpoints**

#### 15.2.1. Specification

Property	Type	Description
<b>addressType</b>	<b>string</b>	<p>addressType specifies the type of address carried by this EndpointSlice. All addresses in this slice must be the same type. This field is immutable after creation. The following address types are currently supported: * IPv4: Represents an IPv4 Address. * IPv6: Represents an IPv6 Address. * FQDN: Represents a Fully Qualified Domain Name.</p> <p>Possible enum values: - <b>"FQDN"</b> represents a FQDN. - <b>"IPv4"</b> represents an IPv4 Address. - <b>"IPv6"</b> represents an IPv6 Address.</p>

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>endpoints</b>	<b>array</b>	endpoints is a list of unique endpoints in this slice. Each slice may include a maximum of 1000 endpoints.
<b>endpoints[]</b>	<b>object</b>	Endpoint represents a single logical "backend" implementing a service.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata.
<b>ports</b>	<b>array</b>	ports specifies the list of network ports exposed by each endpoint in this slice. Each port must have a unique name. When ports is empty, it indicates that there are no defined ports. When a port is defined with a nil port value, it indicates "all ports". Each slice may include a maximum of 100 ports.
<b>ports[]</b>	<b>object</b>	EndpointPort represents a Port used by an EndpointSlice

### 15.2.1.1. .endpoints

#### Description

endpoints is a list of unique endpoints in this slice. Each slice may include a maximum of 1000 endpoints.

#### Type

**array**

### 15.2.1.2. .endpoints[]

#### Description

Endpoint represents a single logical "backend" implementing a service.

#### Type

**object**

#### Required

- **addresses**

Property	Type	Description
<b>addresses</b>	<b>array (string)</b>	addresses of this endpoint. The contents of this field are interpreted according to the corresponding EndpointSlice addressType field. Consumers must handle different types of addresses in the context of their own capabilities. This must contain at least one address but no more than 100. These are all assumed to be fungible and clients may choose to only use the first element. Refer to: <a href="https://issue.k8s.io/106267">https://issue.k8s.io/106267</a>
<b>conditions</b>	<b>object</b>	EndpointConditions represents the current condition of an endpoint.



Property	Type	Description
<b>deprecatedTopology</b>	<b>object (string)</b>	deprecatedTopology contains topology information part of the v1beta1 API. This field is deprecated, and will be removed when the v1beta1 API is removed (no sooner than kubernetes v1.24). While this field can hold values, it is not writable through the v1 API, and any attempts to write to it will be silently ignored. Topology information can be found in the zone and nodeName fields instead.
<b>hints</b>	<b>object</b>	EndpointHints provides hints describing how an endpoint should be consumed.
<b>hostname</b>	<b>string</b>	hostname of this endpoint. This field may be used by consumers of endpoints to distinguish endpoints from each other (e.g. in DNS names). Multiple endpoints which use the same hostname should be considered fungible (e.g. multiple A values in DNS). Must be lowercase and pass DNS Label (RFC 1123) validation.
<b>nodeName</b>	<b>string</b>	nodeName represents the name of the Node hosting this endpoint. This can be used to determine endpoints local to a Node.
<b>targetRef</b>	<b>ObjectReference</b>	targetRef is a reference to a Kubernetes object that represents this endpoint.
<b>zone</b>	<b>string</b>	zone is the name of the Zone this endpoint exists in.

### 15.2.1.3. .endpoints[].conditions

#### Description

EndpointConditions represents the current condition of an endpoint.

#### Type

**object**

Property	Type	Description
<b>ready</b>	<b>boolean</b>	ready indicates that this endpoint is prepared to receive traffic, according to whatever system is managing the endpoint. A nil value indicates an unknown state. In most cases consumers should interpret this unknown state as ready. For compatibility reasons, ready should never be "true" for terminating endpoints, except when the normal readiness behavior is being explicitly overridden, for example when the associated Service has set the publishNotReadyAddresses flag.
<b>serving</b>	<b>boolean</b>	serving is identical to ready except that it is set regardless of the terminating state of endpoints. This condition should be set to true for a ready endpoint that is terminating. If nil, consumers should defer to the ready condition.
<b>terminating</b>	<b>boolean</b>	terminating indicates that this endpoint is terminating. A nil value indicates an unknown state. Consumers should interpret this unknown state to mean that the endpoint is not terminating.

#### 15.2.1.4. .endpoints[].hints

##### Description

EndpointHints provides hints describing how an endpoint should be consumed.

##### Type

**object**

Property	Type	Description
<b>forZones</b>	<b>array</b>	forZones indicates the zone(s) this endpoint should be consumed by to enable topology aware routing.

Property	Type	Description
<b>forZones[]</b>	<b>object</b>	ForZone provides information about which zones should consume this endpoint.

#### 15.2.1.5. .endpoints[].hints.forZones

##### Description

forZones indicates the zone(s) this endpoint should be consumed by to enable topology aware routing.

##### Type

**array**

#### 15.2.1.6. .endpoints[].hints.forZones[]

##### Description

ForZone provides information about which zones should consume this endpoint.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name represents the name of the zone.

#### 15.2.1.7. .ports

##### Description

ports specifies the list of network ports exposed by each endpoint in this slice. Each port must have a unique name. When ports is empty, it indicates that there are no defined ports. When a port is defined with a nil port value, it indicates "all ports". Each slice may include a maximum of 100 ports.

##### Type

**array**

#### 15.2.1.8. .ports[]

##### Description

EndpointPort represents a Port used by an EndpointSlice

##### Type

**object**

Property	Type	Description
<b>appProtocol</b>	<b>string</b>	<p>The application protocol for this port. This is used as a hint for implementations to offer richer behavior for protocols that they understand. This field follows standard Kubernetes label syntax. Valid values are either:</p> <ul style="list-style-type: none"> <li>* Un-prefixed protocol names - reserved for IANA standard service names (as per RFC-6335 and <a href="https://www.iana.org/assignments/service-names">https://www.iana.org/assignments/service-names</a>).</li> <li>* Kubernetes-defined prefixed names: * 'kubernetes.io/h2c' - HTTP/2 over cleartext as described in <a href="https://www.rfc-editor.org/rfc/rfc7540">https://www.rfc-editor.org/rfc/rfc7540</a></li> <li>* Other protocols should use implementation-defined prefixed names such as mycompany.com/my-custom-protocol.</li> </ul>
<b>name</b>	<b>string</b>	<p>name represents the name of this port. All ports in an EndpointSlice must have a unique name. If the EndpointSlice is derived from a Kubernetes service, this corresponds to the Service.ports[].name. Name must either be an empty string or pass DNS_LABEL validation: * must be no more than 63 characters long. * must consist of lower case alphanumeric characters or '-'. * must start and end with an alphanumeric character. Default is empty string.</p>
<b>port</b>	<b>integer</b>	<p>port represents the port number of the endpoint. If this is not specified, ports are not restricted and must be interpreted in the context of the specific consumer.</p>

Property	Type	Description
<b>protocol</b>	<b>string</b>	<p>protocol represents the IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.</p> <p>Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.</p>

### 15.2.2. API endpoints

The following API endpoints are available:

- **/apis/discovery.k8s.io/v1/endpointslices**
  - **GET**: list or watch objects of kind EndpointSlice
- **/apis/discovery.k8s.io/v1/watch/endpointslices**
  - **GET**: watch individual changes to a list of EndpointSlice. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/discovery.k8s.io/v1/namespaces/{namespace}/endpointslices**
  - **DELETE**: delete collection of EndpointSlice
  - **GET**: list or watch objects of kind EndpointSlice
  - **POST**: create an EndpointSlice
- **/apis/discovery.k8s.io/v1/watch/namespaces/{namespace}/endpointslices**
  - **GET**: watch individual changes to a list of EndpointSlice. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/discovery.k8s.io/v1/namespaces/{namespace}/endpointslices/{name}**
  - **DELETE**: delete an EndpointSlice
  - **GET**: read the specified EndpointSlice
  - **PATCH**: partially update the specified EndpointSlice
  - **PUT**: replace the specified EndpointSlice
- **/apis/discovery.k8s.io/v1/watch/namespaces/{namespace}/endpointslices/{name}**
  - **GET**: watch changes to an object of kind EndpointSlice. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

## 15.2.2.1. /apis/discovery.k8s.io/v1/endpointslices

Table 15.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind EndpointSlice

Table 15.2. HTTP responses



HTTP code	Response body
200 - OK	<b>EndpointSliceList</b> schema
401 - Unauthorized	Empty

### 15.2.2.2. /apis/discovery.k8s.io/v1/watch/endpointlices

Table 15.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of EndpointSlice. deprecated: use the 'watch' parameter with a list operation instead.

Table 15.4. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 15.2.2.3. /apis/discovery.k8s.io/v1/namespaces/{namespace}/endpointslices

Table 15.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 15.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of EndpointSlice

Table 15.7. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 15.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 15.9. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind EndpointSlice

Table 15.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 15.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">EndpointSliceList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create an EndpointSlice

**Table 15.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 15.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>EndpointSlice</b> schema	

**Table 15.14. HTTP responses**

HTTP code	Response body
200 - OK	<b>EndpointSlice</b> schema

HTTP code	Response body
201 - Created	<a href="#">EndpointSlice</a> schema
202 - Accepted	<a href="#">EndpointSlice</a> schema
401 - Unauthorized	Empty

#### 15.2.2.4. /apis/discovery.k8s.io/v1/watch/namespaces/{namespace}/endpointslices

Table 15.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 15.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of EndpointSlice. deprecated: use the 'watch' parameter with a list operation instead.

Table 15.17. HTTP responses



HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 15.2.2.5. /apis/discovery.k8s.io/v1/namespaces/{namespace}/endpointslices/{name}

Table 15.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the EndpointSlice
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 15.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete an EndpointSlice

Table 15.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 15.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 15.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified EndpointSlice

Table 15.23. HTTP responses

HTTP code	Response body
200 - OK	<b>EndpointSlice</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified EndpointSlice

**Table 15.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 15.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 15.26. HTTP responses

HTTP code	Response body
200 - OK	<b>EndpointSlice</b> schema
201 - Created	<b>EndpointSlice</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified EndpointSlice

Table 15.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 15.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>EndpointSlice</b> schema	

Table 15.29. HTTP responses

HTTP code	Response body
200 - OK	<b>EndpointSlice</b> schema
201 - Created	<b>EndpointSlice</b> schema
401 - Unauthorized	Empty

### 15.2.2.6. /apis/discovery.k8s.io/v1/watch/namespaces/{namespace}/endpointslices/{name}

Table 15.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the EndpointSlice
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 15.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind EndpointSlice. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 15.32. HTTP responses**



HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## CHAPTER 16. EVENTS APIS

### 16.1. EVENTS APIS

#### 16.1.1. Event [events.k8s.io/v1]

##### Description

Event is a report of an event somewhere in the cluster. It generally denotes some state change in the system. Events have a limited retention time and triggers and messages may evolve with time. Event consumers should not rely on the timing of an event with a given Reason reflecting a consistent underlying trigger, or the continued existence of events with that Reason. Events should be treated as informative, best-effort, supplemental data.

##### Type

**object**

### 16.2. EVENT [EVENTS.K8S.IO/V1]

##### Description

Event is a report of an event somewhere in the cluster. It generally denotes some state change in the system. Events have a limited retention time and triggers and messages may evolve with time. Event consumers should not rely on the timing of an event with a given Reason reflecting a consistent underlying trigger, or the continued existence of events with that Reason. Events should be treated as informative, best-effort, supplemental data.

##### Type

**object**

##### Required

- **eventTime**

#### 16.2.1. Specification

Property	Type	Description
<b>action</b>	<b>string</b>	action is what action was taken/failed regarding to the regarding object. It is machine-readable. This field cannot be empty for new Events and it can have at most 128 characters.

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>deprecatedCount</b>	<b>integer</b>	deprecatedCount is the deprecated field assuring backward compatibility with core.v1 Event type.
<b>deprecatedFirstTimestamp</b>	<b>Time</b>	deprecatedFirstTimestamp is the deprecated field assuring backward compatibility with core.v1 Event type.
<b>deprecatedLastTimestamp</b>	<b>Time</b>	deprecatedLastTimestamp is the deprecated field assuring backward compatibility with core.v1 Event type.
<b>deprecatedSource</b>	<b>EventSource</b>	deprecatedSource is the deprecated field assuring backward compatibility with core.v1 Event type.
<b>eventTime</b>	<b>MicroTime</b>	eventTime is the time when this Event was first observed. It is required.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>note</b>	<b>string</b>	note is a human-readable description of the status of this operation. Maximal length of the note is 1kB, but libraries should be prepared to handle values up to 64kB.
<b>reason</b>	<b>string</b>	reason is why the action was taken. It is human-readable. This field cannot be empty for new Events and it can have at most 128 characters.
<b>regarding</b>	<b>ObjectReference</b>	regarding contains the object this Event is about. In most cases it's an Object reporting controller implements, e.g. ReplicaSetController implements ReplicaSets and this event is emitted because it acts on some changes in a ReplicaSet object.
<b>related</b>	<b>ObjectReference</b>	related is the optional secondary object for more complex actions. E.g. when regarding object triggers a creation or deletion of related object.
<b>reportingController</b>	<b>string</b>	reportingController is the name of the controller that emitted this Event, e.g. <b>kubernetes.io/kubelet</b> . This field cannot be empty for new Events.
<b>reportingInstance</b>	<b>string</b>	reportingInstance is the ID of the controller instance, e.g. <b>kubelet-xyzf</b> . This field cannot be empty for new Events and it can have at most 128 characters.

Property	Type	Description
<b>series</b>	<b>object</b>	EventSeries contain information on series of events, i.e. thing that was/is happening continuously for some time. How often to update the EventSeries is up to the event reporters. The default event reporter in "k8s.io/client-go/tools/events/event_broadcaster.go" shows how this struct is updated on heartbeats and can guide customized reporter implementations.
<b>type</b>	<b>string</b>	type is the type of this event (Normal, Warning), new types could be added in the future. It is machine-readable. This field cannot be empty for new Events.

### 16.2.1.1. .series

#### Description

EventSeries contain information on series of events, i.e. thing that was/is happening continuously for some time. How often to update the EventSeries is up to the event reporters. The default event reporter in "k8s.io/client-go/tools/events/event\_broadcaster.go" shows how this struct is updated on heartbeats and can guide customized reporter implementations.

#### Type

**object**

#### Required

- **count**
- **lastObservedTime**

Property	Type	Description
<b>count</b>	<b>integer</b>	count is the number of occurrences in this series up to the last heartbeat time.
<b>lastObservedTime</b>	<b>MicroTime</b>	lastObservedTime is the time when last Event from the series was seen before last heartbeat.

### 16.2.2. API endpoints

The following API endpoints are available:

- **/apis/events.k8s.io/v1/events**
  - **GET**: list or watch objects of kind Event
- **/apis/events.k8s.io/v1/watch/events**
  - **GET**: watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/events.k8s.io/v1/namespaces/{namespace}/events**
  - **DELETE**: delete collection of Event
  - **GET**: list or watch objects of kind Event
  - **POST**: create an Event
- **/apis/events.k8s.io/v1/watch/namespaces/{namespace}/events**
  - **GET**: watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/events.k8s.io/v1/namespaces/{namespace}/events/{name}**
  - **DELETE**: delete an Event
  - **GET**: read the specified Event
  - **PATCH**: partially update the specified Event
  - **PUT**: replace the specified Event
- **/apis/events.k8s.io/v1/watch/namespaces/{namespace}/events/{name}**
  - **GET**: watch changes to an object of kind Event. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 16.2.2.1. /apis/events.k8s.io/v1/events

Table 16.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Event

**Table 16.2. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>EventList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 16.2.2.2. /apis/events.k8s.io/v1/watch/events

Table 16.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.

**Table 16.4. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 16.2.2.3. /apis/events.k8s.io/v1/namespaces/{namespace}/events

Table 16.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 16.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Event

Table 16.7. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>



Table 16.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 16.9. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Event

Table 16.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 16.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">EventList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create an Event

**Table 16.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 16.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Event</b> schema	

**Table 16.14. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Event</b> schema

HTTP code	Response body
201 - Created	<b>Event</b> schema
202 - Accepted	<b>Event</b> schema
401 - Unauthorized	Empty

#### 16.2.2.4. /apis/events.k8s.io/v1/watch/namespaces/{namespace}/events

Table 16.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 16.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Event. deprecated: use the 'watch' parameter with a list operation instead.

Table 16.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 16.2.2.5. /apis/events.k8s.io/v1/namespaces/{namespace}/events/{name}

Table 16.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Event
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 16.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete an Event

Table 16.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 16.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 16.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Event

Table 16.23. HTTP responses

HTTP code	Response body
200 - OK	<b>Event</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Event

**Table 16.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 16.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 16.26. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Event</b> schema
201 - Created	<b>Event</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Event

Table 16.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 16.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Event</b> schema	

Table 16.29. HTTP responses

HTTP code	Response body
200 - OK	<b>Event</b> schema
201 - Created	<b>Event</b> schema
401 - Unauthorized	Empty

#### 16.2.2.6. /apis/events.k8s.io/v1/watch/namespaces/{namespace}/events/{name}

Table 16.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Event
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 16.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Event. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 16.32. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## CHAPTER 17. FLOW CONTROL APIS

### 17.1. FLOW CONTROL APIS

#### 17.1.1. FlowSchema [flowcontrol.apiserver.k8s.io/v1beta3]

##### Description

FlowSchema defines the schema of a group of flows. Note that a flow is made up of a set of inbound API requests with similar attributes and is identified by a pair of strings: the name of the FlowSchema and a "flow distinguisher".

##### Type

**object**

#### 17.1.2. PriorityLevelConfiguration [flowcontrol.apiserver.k8s.io/v1beta3]

##### Description

PriorityLevelConfiguration represents the configuration of a priority level.

##### Type

**object**

### 17.2. FLOWSHEMA [FLOWCONTROL.APISERVER.K8S.IO/V1BETA3]

##### Description

FlowSchema defines the schema of a group of flows. Note that a flow is made up of a set of inbound API requests with similar attributes and is identified by a pair of strings: the name of the FlowSchema and a "flow distinguisher".

##### Type

**object**

#### 17.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	<b>metadata</b> is the standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	FlowSchemaSpec describes how the FlowSchema's specification looks like.
<b>status</b>	<b>object</b>	FlowSchemaStatus represents the current state of a FlowSchema.

### 17.2.1.1. .spec

#### Description

FlowSchemaSpec describes how the FlowSchema's specification looks like.

#### Type

**object**

#### Required

- **priorityLevelConfiguration**

Property	Type	Description
<b>distinguisherMethod</b>	<b>object</b>	FlowDistinguisherMethod specifies the method of a flow distinguisher.

Property	Type	Description
<b>matchingPrecedence</b>	<b>integer</b>	<b>matchingPrecedence</b> is used to choose among the FlowSchemas that match a given request. The chosen FlowSchema is among those with the numerically lowest (which we take to be logically highest) MatchingPrecedence. Each MatchingPrecedence value must be ranged in [1,10000]. Note that if the precedence is not specified, it will be set to 1000 as default.
<b>priorityLevelConfiguration</b>	<b>object</b>	PriorityLevelConfigurationReference contains information that points to the "request-priority" being used.
<b>rules</b>	<b>array</b>	<b>rules</b> describes which requests will match this flow schema. This FlowSchema matches a request if and only if at least one member of rules matches the request. if it is an empty slice, there will be no requests matching the FlowSchema.
<b>rules[]</b>	<b>object</b>	PolicyRulesWithSubjects prescribes a test that applies to a request to an apiserver. The test considers the subject making the request, the verb being requested, and the resource to be acted upon. This PolicyRulesWithSubjects matches a request if and only if both (a) at least one member of subjects matches the request and (b) at least one member of resourceRules or nonResourceRules matches the request.

### 17.2.1.2. .spec.distinguisherMethod

#### Description

FlowDistinguisherMethod specifies the method of a flow distinguisher.

#### Type

**object**

**Required**

- **type**

Property	Type	Description
<b>type</b>	<b>string</b>	<b>type</b> is the type of flow distinguisher method The supported types are "ByUser" and "ByNamespace". Required.

**17.2.1.3. .spec.priorityLevelConfiguration****Description**

PriorityLevelConfigurationReference contains information that points to the "request-priority" being used.

**Type**

**object**

**Required**

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	<b>name</b> is the name of the priority level configuration being referenced Required.

**17.2.1.4. .spec.rules****Description**

**rules** describes which requests will match this flow schema. This FlowSchema matches a request if and only if at least one member of rules matches the request. if it is an empty slice, there will be no requests matching the FlowSchema.

**Type**

**array**

**17.2.1.5. .spec.rules[]****Description**

PolicyRulesWithSubjects prescribes a test that applies to a request to an apiserver. The test considers the subject making the request, the verb being requested, and the resource to be acted upon. This PolicyRulesWithSubjects matches a request if and only if both (a) at least one member of subjects matches the request and (b) at least one member of resourceRules or nonResourceRules matches the request.

**Type**

**object**

## Required

- **subjects**

Property	Type	Description
<b>nonResourceRules</b>	<b>array</b>	<b>nonResourceRules</b> is a list of NonResourcePolicyRules that identify matching requests according to their verb and the target non-resource URL.
<b>nonResourceRules[]</b>	<b>object</b>	NonResourcePolicyRule is a predicate that matches non-resource requests according to their verb and the target non-resource URL. A NonResourcePolicyRule matches a request if and only if both (a) at least one member of verbs matches the request and (b) at least one member of nonResourceURLs matches the request.
<b>resourceRules</b>	<b>array</b>	<b>resourceRules</b> is a slice of ResourcePolicyRules that identify matching requests according to their verb and the target resource. At least one of <b>resourceRules</b> and <b>nonResourceRules</b> has to be non-empty.

Property	Type	Description
<b>resourceRules[]</b>	<b>object</b>	ResourcePolicyRule is a predicate that matches some resource requests, testing the request's verb and the target resource. A ResourcePolicyRule matches a resource request if and only if: (a) at least one member of verbs matches the request, (b) at least one member of apiGroups matches the request, (c) at least one member of resources matches the request, and (d) either (d1) the request does not specify a namespace (i.e., <b>Namespace=""</b> ) and clusterScope is true or (d2) the request specifies a namespace and least one member of namespaces matches the request's namespace.
<b>subjects</b>	<b>array</b>	subjects is the list of normal user, serviceaccount, or group that this rule cares about. There must be at least one member in this slice. A slice that includes both the system:authenticated and system:unauthenticated user groups matches every request. Required.
<b>subjects[]</b>	<b>object</b>	Subject matches the originator of a request, as identified by the request authentication system. There are three ways of matching an originator; by user, group, or service account.

#### 17.2.1.6. .spec.rules[].nonResourceRules

##### Description

**nonResourceRules** is a list of NonResourcePolicyRules that identify matching requests according to their verb and the target non-resource URL.

##### Type

**array**

#### 17.2.1.7. .spec.rules[].nonResourceRules[]



## Description

NonResourcePolicyRule is a predicate that matches non-resource requests according to their verb and the target non-resource URL. A NonResourcePolicyRule matches a request if and only if both (a) at least one member of verbs matches the request and (b) at least one member of nonResourceURLs matches the request.

## Type

**object**

## Required

- **verbs**
- **nonResourceURLs**

Property	Type	Description
<b>nonResourceURLs</b>	<b>array (string)</b>	<b>nonResourceURLs</b> is a set of url prefixes that a user should have access to and may not be empty. For example: - "/healthz" is legal - "/hea*" is illegal - "/hea" is legal but matches nothing - "/hea/" <b>also matches nothing</b> - <b>"/healthz/"</b> matches all per-component health checks. "*" matches all non-resource urls. if it is present, it must be the only entry. Required.
<b>verbs</b>	<b>array (string)</b>	<b>verbs</b> is a list of matching verbs and may not be empty. "*" matches all verbs. If it is present, it must be the only entry. Required.

### 17.2.1.8. .spec.rules[].resourceRules

## Description

**resourceRules** is a slice of ResourcePolicyRules that identify matching requests according to their verb and the target resource. At least one of **resourceRules** and **nonResourceRules** has to be non-empty.

## Type

**array**

### 17.2.1.9. .spec.rules[].resourceRules[]

## Description

ResourcePolicyRule is a predicate that matches some resource requests, testing the request's verb and the target resource. A ResourcePolicyRule matches a resource request if and only if: (a) at least one member of verbs matches the request, (b) at least one member of apiGroups matches the

request, (c) at least one member of resources matches the request, and (d) either (d1) the request does not specify a namespace (i.e., **Namespace==""**) and **clusterScope** is true or (d2) the request specifies a namespace and least one member of namespaces matches the request's namespace.

## Type

**object**

## Required

- **verbs**
- **apiGroups**
- **resources**

Property	Type	Description
<b>apiGroups</b>	<b>array (string)</b>	<b>apiGroups</b> is a list of matching API groups and may not be empty. "*" matches all API groups and, if present, must be the only entry. Required.
<b>clusterScope</b>	<b>boolean</b>	<b>clusterScope</b> indicates whether to match requests that do not specify a namespace (which happens either because the resource is not namespaced or the request targets all namespaces). If this field is omitted or false then the <b>namespaces</b> field must contain a non-empty list.
<b>namespaces</b>	<b>array (string)</b>	<b>namespaces</b> is a list of target namespaces that restricts matches. A request that specifies a target namespace matches only if either (a) this list contains that target namespace or (b) this list contains "". <b>Note that ""</b> matches any specified namespace but does not match a request that <i>does not specify</i> a namespace (see the <b>clusterScope</b> field for that). This list may be empty, but only if <b>clusterScope</b> is true.

Property	Type	Description
<b>resources</b>	<b>array (string)</b>	<b>resources</b> is a list of matching resources (i.e., lowercase and plural) with, if desired, subresource. For example, [ "services", "nodes/status" ]. This list may not be empty. "*" matches all resources and, if present, must be the only entry. Required.
<b>verbs</b>	<b>array (string)</b>	<b>verbs</b> is a list of matching verbs and may not be empty. "*" matches all verbs and, if present, must be the only entry. Required.

#### 17.2.1.10. .spec.rules[].subjects

##### Description

subjects is the list of normal user, serviceaccount, or group that this rule cares about. There must be at least one member in this slice. A slice that includes both the system:authenticated and system:unauthenticated user groups matches every request. Required.

##### Type

**array**

#### 17.2.1.11. .spec.rules[].subjects[]

##### Description

Subject matches the originator of a request, as identified by the request authentication system. There are three ways of matching an originator; by user, group, or service account.

##### Type

**object**

##### Required

- **kind**

Property	Type	Description
<b>group</b>	<b>object</b>	GroupSubject holds detailed information for group-kind subject.
<b>kind</b>	<b>string</b>	<b>kind</b> indicates which one of the other fields is non-empty. Required

Property	Type	Description
<b>serviceAccount</b>	<b>object</b>	ServiceAccountSubject holds detailed information for service-account-kind subject.
<b>user</b>	<b>object</b>	UserSubject holds detailed information for user-kind subject.

#### 17.2.1.12. .spec.rules[].subjects[].group

##### Description

GroupSubject holds detailed information for group-kind subject.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the user group that matches, or "*" to match all user groups. See <a href="https://github.com/kubernetes/apiserver/blob/master/pkg/authentication/user/user.go">https://github.com/kubernetes/apiserver/blob/master/pkg/authentication/user/user.go</a> for some well-known group names. Required.

#### 17.2.1.13. .spec.rules[].subjects[].serviceAccount

##### Description

ServiceAccountSubject holds detailed information for service-account-kind subject.

##### Type

**object**

##### Required

- **namespace**
- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	<b>name</b> is the name of matching ServiceAccount objects, or "*" to match regardless of name. Required.
<b>namespace</b>	<b>string</b>	<b>namespace</b> is the namespace of matching ServiceAccount objects. Required.

#### 17.2.1.14. .spec.rules[].subjects[].user

##### Description

UserSubject holds detailed information for user-kind subject.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	<b>name</b> is the username that matches, or "*" to match all usernames. Required.

#### 17.2.1.15. .status

##### Description

FlowSchemaStatus represents the current state of a FlowSchema.

##### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array</b>	<b>conditions</b> is a list of the current states of FlowSchema.
<b>conditions[]</b>	<b>object</b>	FlowSchemaCondition describes conditions for a FlowSchema.

#### 17.2.1.16. .status.conditions

##### Description

**conditions** is a list of the current states of FlowSchema.

Type

**array**

### 17.2.1.17. .status.conditions[]

Description

FlowSchemaCondition describes conditions for a FlowSchema.

Type

**object**

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	<b>lastTransitionTime</b> is the last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	<b>message</b> is a human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	<b>reason</b> is a unique, one-word, CamelCase reason for the condition's last transition.
<b>status</b>	<b>string</b>	<b>status</b> is the status of the condition. Can be True, False, Unknown. Required.
<b>type</b>	<b>string</b>	<b>type</b> is the type of the condition. Required.

## 17.2.2. API endpoints

The following API endpoints are available:

- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas**
  - **DELETE**: delete collection of FlowSchema
  - **GET**: list or watch objects of kind FlowSchema
  - **POST**: create a FlowSchema
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/flowschemas**
  - **GET**: watch individual changes to a list of FlowSchema. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas/{name}**
  - **DELETE**: delete a FlowSchema

- **GET**: read the specified FlowSchema
- **PATCH**: partially update the specified FlowSchema
- **PUT**: replace the specified FlowSchema
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/flowschemas/{name}**
  - **GET**: watch changes to an object of kind FlowSchema. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas/{name}/status**
  - **GET**: read status of the specified FlowSchema
  - **PATCH**: partially update status of the specified FlowSchema
  - **PUT**: replace status of the specified FlowSchema

### 17.2.2.1. /apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas

Table 17.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of FlowSchema

Table 17.2. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>`continue`</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 17.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 17.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind FlowSchema

Table 17.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 17.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">FlowSchemaList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a FlowSchema

**Table 17.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 17.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>FlowSchema</b> schema	

**Table 17.9. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>FlowSchema</b> schema

HTTP code	Response body
201 - Created	<b>FlowSchema</b> schema
202 - Accepted	<b>FlowSchema</b> schema
401 - Unauthorized	Empty

### 17.2.2.2. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/flowschemas

Table 17.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>



Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
**GET**

**Description**

watch individual changes to a list of FlowSchema. deprecated: use the 'watch' parameter with a list operation instead.

**Table 17.11. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**17.2.2.3. /apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas/{name}****Table 17.12. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the FlowSchema

**Table 17.13. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a FlowSchema

**Table 17.14. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 17.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 17.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified FlowSchema

Table 17.17. HTTP responses

HTTP code	Response body
200 - OK	<b>FlowSchema</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified FlowSchema

**Table 17.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 17.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 17.20. HTTP responses

HTTP code	Response body
200 - OK	<b>FlowSchema</b> schema
201 - Created	<b>FlowSchema</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified FlowSchema

Table 17.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 17.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>FlowSchema</b> schema	

Table 17.23. HTTP responses

HTTP code	Response body
200 - OK	<b>FlowSchema</b> schema
201 - Created	<b>FlowSchema</b> schema
401 - Unauthorized	Empty

#### 17.2.2.4. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/flowschemas/{name}

Table 17.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the FlowSchema

Table 17.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind FlowSchema. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 17.26. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 17.2.2.5. /apis/flowcontrol.apiserver.k8s.io/v1beta3/flowschemas/{name}/status

Table 17.27. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the FlowSchema

Table 17.28. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified FlowSchema

Table 17.29. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">FlowSchema</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified FlowSchema

Table 17.30. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 17.31. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 17.32. HTTP responses

HTTP code	Response body
200 - OK	<b>FlowSchema</b> schema
201 - Created	<b>FlowSchema</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified FlowSchema

**Table 17.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 17.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>FlowSchema</b> schema	

Table 17.35. HTTP responses

HTTP code	Response body
200 - OK	<b>FlowSchema</b> schema
201 - Created	<b>FlowSchema</b> schema
401 - Unauthorized	Empty

## 17.3. PRIORITYLEVELCONFIGURATION [FLOWCONTROL.APISERVER.K8S.IO/V1BETA3]

### Description

PriorityLevelConfiguration represents the configuration of a priority level.

### Type

**object**

### 17.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	<b>metadata</b> is the standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	PriorityLevelConfigurationSpec specifies the configuration of a priority level.
<b>status</b>	<b>object</b>	PriorityLevelConfigurationStatus represents the current state of a "request-priority".

### 17.3.1.1. .spec

#### Description

PriorityLevelConfigurationSpec specifies the configuration of a priority level.

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>limited</b>	<b>object</b>	LimitedPriorityLevelConfiguration specifies how to handle requests that are subject to limits. It addresses two issues: - How are requests for this priority level limited? - What should be done with requests that exceed the limit?

Property	Type	Description
<b>type</b>	<b>string</b>	<b>type</b> indicates whether this priority level is subject to limitation on request execution. A value of " <b>Exempt</b> " means that requests of this priority level are not subject to a limit (and thus are never queued) and do not detract from the capacity made available to other priority levels. A value of " <b>Limited</b> " means that (a) requests of this priority level are subject to limits and (b) some of the server's limited capacity is made available exclusively to this priority level. Required.

### 17.3.1.2. .spec.limited

#### Description

LimitedPriorityLevelConfiguration specifies how to handle requests that are subject to limits. It addresses two issues: - How are requests for this priority level limited? - What should be done with requests that exceed the limit?

#### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>borrowingLimitPercent</b>	<b>integer</b>	<p><b>borrowingLimitPercent</b>, if present, configures a limit on how many seats this priority level can borrow from other priority levels. The limit is known as this level's BorrowingConcurrencyLimit (BorrowingCL) and is a limit on the total number of seats that this level may borrow at any one time. This field holds the ratio of that limit to the level's nominal concurrency limit. When this field is non-nil, it must hold a non-negative integer and the limit is calculated as follows.</p> $\text{BorrowingCL}(i) = \text{round}(\text{NominalCL}(i) * \text{borrowingLimitPercent}(i)/100.0)$ <p>The value of this field can be more than 100, implying that this priority level can borrow a number of seats that is greater than its own nominal concurrency limit (NominalCL). When this field is left <b>nil</b>, the limit is effectively infinite.</p>
<b>lendablePercent</b>	<b>integer</b>	<p><b>lendablePercent</b> prescribes the fraction of the level's NominalCL that can be borrowed by other priority levels. The value of this field must be between 0 and 100, inclusive, and it defaults to 0. The number of seats that other levels can borrow from this level, known as this level's LendableConcurrencyLimit (LendableCL), is defined as follows.</p> $\text{LendableCL}(i) = \text{round}(\text{NominalCL}(i) * \text{lendablePercent}(i)/100.0)$
<b>limitResponse</b>	<b>object</b>	LimitResponse defines how to handle requests that can not be executed right now.

Property	Type	Description
<b>nominalConcurrencyShares</b>	<b>integer</b>	<p><b>nominalConcurrencyShares</b> (NCS) contributes to the computation of the <b>NominalConcurrencyLimit</b> (NominalCL) of this level. This is the number of execution seats available at this priority level. This is used both for requests dispatched from this priority level as well as requests dispatched from other priority levels borrowing seats from this level. The server's concurrency limit (ServerCL) is divided among the Limited priority levels in proportion to their NCS values:</p> $\text{NominalCL}(i) = \text{ceil}(\text{ServerCL} * \text{NCS}(i) / \text{sum\_ncs})$ $\text{sum\_ncs} = \text{sum}[\text{limited priority level } k] \text{NCS}(k)$ <p>Bigger numbers mean a larger nominal concurrency limit, at the expense of every other Limited priority level. This field has a default value of 30.</p>

### 17.3.1.3. .spec.limited.limitResponse

#### Description

LimitResponse defines how to handle requests that can not be executed right now.

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>queuing</b>	<b>object</b>	QueuingConfiguration holds the configuration parameters for queuing

Property	Type	Description
<b>type</b>	<b>string</b>	<b>type</b> is "Queue" or "Reject". "Queue" means that requests that can not be executed upon arrival are held in a queue until they can be executed or a queuing limit is reached. "Reject" means that requests that can not be executed upon arrival are rejected. Required.

#### 17.3.1.4. .spec.limited.limitResponse.queuing

##### Description

QueuingConfiguration holds the configuration parameters for queuing

##### Type

**object**

Property	Type	Description
<b>handSize</b>	<b>integer</b>	<b>handSize</b> is a small positive number that configures the shuffle sharding of requests into queues. When enqueueing a request at this priority level the request's flow identifier (a string pair) is hashed and the hash value is used to shuffle the list of queues and deal a hand of the size specified here. The request is put into one of the shortest queues in that hand. <b>handSize</b> must be no larger than <b>queues</b> , and should be significantly smaller (so that a few heavy flows do not saturate most of the queues). See the user-facing documentation for more extensive guidance on setting this field. This field has a default value of 8.
<b>queueLengthLimit</b>	<b>integer</b>	<b>queueLengthLimit</b> is the maximum number of requests allowed to be waiting in a given queue of this priority level at a time; excess requests are rejected. This value must be positive. If not specified, it will be defaulted to 50.

Property	Type	Description
<b>queues</b>	<b>integer</b>	<b>queues</b> is the number of queues for this priority level. The queues exist independently at each apiserver. The value must be positive. Setting it to 1 effectively precludes shuffleharding and thus makes the distinguisher method of associated flow schemas irrelevant. This field has a default value of 64.

### 17.3.1.5. .status

#### Description

PriorityLevelConfigurationStatus represents the current state of a "request-priority".

#### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array</b>	<b>conditions</b> is the current state of "request-priority".
<b>conditions[]</b>	<b>object</b>	PriorityLevelConfigurationCondition defines the condition of priority level.

### 17.3.1.6. .status.conditions

#### Description

**conditions** is the current state of "request-priority".

#### Type

**array**

### 17.3.1.7. .status.conditions[]

#### Description

PriorityLevelConfigurationCondition defines the condition of priority level.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>lastTransitionTime</b>	<b>Time</b>	<b>lastTransitionTime</b> is the last time the condition transitioned from one status to another.
<b>message</b>	<b>string</b>	<b>message</b> is a human-readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	<b>reason</b> is a unique, one-word, CamelCase reason for the condition's last transition.
<b>status</b>	<b>string</b>	<b>status</b> is the status of the condition. Can be True, False, Unknown. Required.
<b>type</b>	<b>string</b>	<b>type</b> is the type of the condition. Required.

### 17.3.2. API endpoints

The following API endpoints are available:

- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations**
  - **DELETE**: delete collection of PriorityLevelConfiguration
  - **GET**: list or watch objects of kind PriorityLevelConfiguration
  - **POST**: create a PriorityLevelConfiguration
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/prioritylevelconfigurations**
  - **GET**: watch individual changes to a list of PriorityLevelConfiguration. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations/{name}**
  - **DELETE**: delete a PriorityLevelConfiguration
  - **GET**: read the specified PriorityLevelConfiguration
  - **PATCH**: partially update the specified PriorityLevelConfiguration
  - **PUT**: replace the specified PriorityLevelConfiguration
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/prioritylevelconfigurations/{name}**

- **GET**: watch changes to an object of kind PriorityLevelConfiguration. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations/{name}/status**
  - **GET**: read status of the specified PriorityLevelConfiguration
  - **PATCH**: partially update status of the specified PriorityLevelConfiguration
  - **PUT**: replace status of the specified PriorityLevelConfiguration

### 17.3.2.1. /apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations

Table 17.36. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of PriorityLevelConfiguration

Table 17.37. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>

Parameter	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 17.38. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 17.39. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind PriorityLevelConfiguration

Table 17.40. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 17.41. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PriorityLevelConfigurationList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a PriorityLevelConfiguration

**Table 17.42. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 17.43. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PriorityLevelConfiguration</b> schema	

**Table 17.44. HTTP responses**

HTTP code	Response body
200 - OK	<b>PriorityLevelConfiguration</b> schema

HTTP code	Response body
201 - Created	<a href="#">PriorityLevelConfiguration</a> schema
202 - Accepted	<a href="#">PriorityLevelConfiguration</a> schema
401 - Unauthorized	Empty

### 17.3.2.2. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/prioritylevelconfigurations

Table 17.45. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PriorityLevelConfiguration. deprecated: use the 'watch' parameter with a list operation instead.

Table 17.46. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 17.3.2.3. /apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations/{name}

Table 17.47. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PriorityLevelConfiguration

Table 17.48. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a PriorityLevelConfiguration

Table 17.49. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 17.50. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 17.51. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified PriorityLevelConfiguration

Table 17.52. HTTP responses

HTTP code	Response body
200 - OK	<b>PriorityLevelConfiguration</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified PriorityLevelConfiguration

**Table 17.53. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 17.54. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 17.55. HTTP responses

HTTP code	Response body
200 - OK	<b>PriorityLevelConfiguration</b> schema
201 - Created	<b>PriorityLevelConfiguration</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified PriorityLevelConfiguration

Table 17.56. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 17.57. Body parameters

Parameter	Type	Description
<b>body</b>	<b>PriorityLevelConfiguration</b> schema	

Table 17.58. HTTP responses

HTTP code	Response body
200 - OK	<b>PriorityLevelConfiguration</b> schema
201 - Created	<b>PriorityLevelConfiguration</b> schema
401 - Unauthorized	Empty

#### 17.3.2.4. /apis/flowcontrol.apiserver.k8s.io/v1beta3/watch/prioritylevelconfigurations/{name}

Table 17.59. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PriorityLevelConfiguration

Table 17.60. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind PriorityLevelConfiguration. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 17.61. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 17.3.2.5. /apis/flowcontrol.apiserver.k8s.io/v1beta3/prioritylevelconfigurations/{name}/status

Table 17.62. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PriorityLevelConfiguration

Table 17.63. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified PriorityLevelConfiguration

Table 17.64. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PriorityLevelConfiguration</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified PriorityLevelConfiguration

Table 17.65. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 17.66. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 17.67. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PriorityLevelConfiguration</a> schema
201 - Created	<a href="#">PriorityLevelConfiguration</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified [PriorityLevelConfiguration](#)

**Table 17.68. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 17.69. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PriorityLevelConfiguration</b> schema	

Table 17.70. HTTP responses

HTTP code	Reponse body
200 - OK	<b>PriorityLevelConfiguration</b> schema
201 - Created	<b>PriorityLevelConfiguration</b> schema
401 - Unauthorized	Empty

## CHAPTER 18. NETWORKING APIS

### 18.1. NETWORKING APIS

#### 18.1.1. Ingress [networking.k8s.io/v1]

##### Description

Ingress is a collection of rules that allow inbound connections to reach the endpoints defined by a backend. An Ingress can be configured to give services externally-reachable urls, load balance traffic, terminate SSL, offer name based virtual hosting etc.

##### Type

**object**

#### 18.1.2. IngressClass [networking.k8s.io/v1]

##### Description

IngressClass represents the class of the Ingress, referenced by the Ingress Spec. The **ingressclass.kubernetes.io/is-default-class** annotation can be used to indicate that an IngressClass should be considered default. When a single IngressClass resource has this annotation set to true, new Ingress resources without a class specified will be assigned this default class.

##### Type

**object**

#### 18.1.3. NetworkPolicy [networking.k8s.io/v1]

##### Description

NetworkPolicy describes what network traffic is allowed for a set of Pods

##### Type

**object**

### 18.2. INGRESS [NETWORKING.K8S.IO/V1]

##### Description

Ingress is a collection of rules that allow inbound connections to reach the endpoints defined by a backend. An Ingress can be configured to give services externally-reachable urls, load balance traffic, terminate SSL, offer name based virtual hosting etc.

##### Type

**object**

#### 18.2.1. Specification

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	IngressSpec describes the Ingress the user wishes to exist.
<b>status</b>	<b>object</b>	IngressStatus describe the current state of the Ingress.

### 18.2.1.1. .spec

#### Description

IngressSpec describes the Ingress the user wishes to exist.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>defaultBackend</b>	<b>object</b>	IngressBackend describes all endpoints for a given service and port.
<b>ingressClassName</b>	<b>string</b>	ingressClassName is the name of an IngressClass cluster resource. Ingress controller implementations use this field to know whether they should be serving this Ingress resource, by a transitive connection (controller → IngressClass → Ingress resource). Although the <b>kubernetes.io/ingress.class</b> annotation (simple constant name) was never formally defined, it was widely supported by Ingress controllers to create a direct binding between Ingress controller and Ingress resources. Newly created Ingress resources should prefer using the field. However, even though the annotation is officially deprecated, for backwards compatibility reasons, ingress controllers should still honor that annotation if present.
<b>rules</b>	<b>array</b>	rules is a list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
<b>rules[]</b>	<b>object</b>	IngressRule represents the rules mapping the paths under a specified host to the related backend services. Incoming requests are first evaluated for a host match, then routed to the backend associated with the matching IngressRuleValue.



Property	Type	Description
<b>tls</b>	<b>array</b>	tls represents the TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.
<b>tls[]</b>	<b>object</b>	IngressTLS describes the transport layer security associated with an ingress.

### 18.2.1.2. .spec.defaultBackend

#### Description

IngressBackend describes all endpoints for a given service and port.

#### Type

**object**

Property	Type	Description
<b>resource</b>	<b>TypedLocalObjectReference</b>	resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, a service.Name and service.Port must not be specified. This is a mutually exclusive setting with "Service".
<b>service</b>	<b>object</b>	IngressServiceBackend references a Kubernetes Service as a Backend.

### 18.2.1.3. .spec.defaultBackend.service

#### Description

IngressServiceBackend references a Kubernetes Service as a Backend.

#### Type

**object**

#### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the referenced service. The service must exist in the same namespace as the Ingress object.
<b>port</b>	<b>object</b>	ServiceBackendPort is the service port being referenced.

#### 18.2.1.4. .spec.defaultBackend.service.port

##### Description

ServiceBackendPort is the service port being referenced.

##### Type

**object**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the port on the Service. This is a mutually exclusive setting with "Number".
<b>number</b>	<b>integer</b>	number is the numerical port number (e.g. 80) on the Service. This is a mutually exclusive setting with "Name".

#### 18.2.1.5. .spec.rules

##### Description

rules is a list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.

##### Type

**array**

#### 18.2.1.6. .spec.rules[]

##### Description

IngressRule represents the rules mapping the paths under a specified host to the related backend services. Incoming requests are first evaluated for a host match, then routed to the backend associated with the matching IngressRuleValue.

##### Type

**object**

Property	Type	Description
<b>host</b>	<b>string</b>	<p>host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <b>The wildcard character "</b> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If host is precise, the request matches this rule if the http host header is equal to Host. 2. If host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>

Property	Type	Description
<b>http</b>	<b>object</b>	HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: <a href="http://&lt;host&gt;/&lt;path&gt;?&lt;searchpart&gt;">http://&lt;host&gt;/&lt;path&gt;?&lt;searchpart&gt;</a> → backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.

### 18.2.1.7. .spec.rules[].http

#### Description

HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: <http://<host>/<path>?<searchpart>> → backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.

#### Type

**object**

#### Required

- **paths**

Property	Type	Description
<b>paths</b>	<b>array</b>	paths is a collection of paths that map requests to backends.
<b>paths[]</b>	<b>object</b>	HTTPIngressPath associates a path with a backend. Incoming urls matching the path are forwarded to the backend.

### 18.2.1.8. .spec.rules[].http.paths

#### Description

paths is a collection of paths that map requests to backends.

#### Type

**array**

### 18.2.1.9. .spec.rules[].http.paths[]

#### Description

HTTPIngressPath associates a path with a backend. Incoming urls matching the path are forwarded to the backend.

Type

**object**

Required

- **pathType**
- **backend**

Property	Type	Description
<b>backend</b>	<b>object</b>	IngressBackend describes all endpoints for a given service and port.
<b>path</b>	<b>string</b>	path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/' and must be present when using PathType with value "Exact" or "Prefix".
<b>pathType</b>	<b>string</b>	pathType determines the interpretation of the path matching. PathType can be one of the following values: <ul style="list-style-type: none"> <li>* Exact: Matches the URL path exactly.</li> <li>* Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz).</li> <li>* ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types.</li> </ul>

Property	Type	Description
		<p>Possible enum values: - <b>"Exact"</b> matches the URL path exactly and with case sensitivity. - <b>"ImplementationSpecific"</b> matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. - <b>"Prefix"</b> matches based on a URL path prefix split by '/'. Matching is case sensitive and done on a path element by element basis. A path element refers to the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). If multiple matching paths exist in an Ingress spec, the longest matching path is given priority. Examples: - /foo/bar does not match requests to /foo/barbaz - /foo/bar matches request to /foo/bar and /foo/bar/baz - /foo and /foo/ both match requests to /foo and /foo/. If both paths are present in an Ingress spec, the longest matching path (/foo/) is given priority.</p>

### 18.2.1.10. .spec.rules[].http.paths[].backend

#### Description

IngressBackend describes all endpoints for a given service and port.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>resource</b>	<b>TypedLocalObjectReference</b>	resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, a service.Name and service.Port must not be specified. This is a mutually exclusive setting with "Service".
<b>service</b>	<b>object</b>	IngressServiceBackend references a Kubernetes Service as a Backend.

#### 18.2.1.11. .spec.rules[].http.paths[].backend.service

##### Description

IngressServiceBackend references a Kubernetes Service as a Backend.

##### Type

**object**

##### Required

- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	name is the referenced service. The service must exist in the same namespace as the Ingress object.
<b>port</b>	<b>object</b>	ServiceBackendPort is the service port being referenced.

#### 18.2.1.12. .spec.rules[].http.paths[].backend.service.port

##### Description

ServiceBackendPort is the service port being referenced.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>name</b>	<b>string</b>	name is the name of the port on the Service. This is a mutually exclusive setting with "Number".
<b>number</b>	<b>integer</b>	number is the numerical port number (e.g. 80) on the Service. This is a mutually exclusive setting with "Name".

### 18.2.1.13. .spec.tls

#### Description

tls represents the TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

#### Type

**array**

### 18.2.1.14. .spec.tls[]

#### Description

IngressTLS describes the transport layer security associated with an ingress.

#### Type

**object**

Property	Type	Description
<b>hosts</b>	<b>array (string)</b>	hosts is a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
<b>secretName</b>	<b>string</b>	secretName is the name of the secret used to terminate TLS traffic on port 443. Field is left optional to allow TLS routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the "Host" header is used for routing.



### 18.2.1.15. `.status`

#### Description

`IngressStatus` describe the current state of the Ingress.

#### Type

**object**

Property	Type	Description
<b>loadBalancer</b>	<b>object</b>	<code>IngressLoadBalancerStatus</code> represents the status of a load-balancer.

### 18.2.1.16. `.status.loadBalancer`

#### Description

`IngressLoadBalancerStatus` represents the status of a load-balancer.

#### Type

**object**

Property	Type	Description
<b>ingress</b>	<b>array</b>	<code>ingress</code> is a list containing ingress points for the load-balancer.
<b>ingress[]</b>	<b>object</b>	<code>IngressLoadBalancerIngress</code> represents the status of a load-balancer ingress point.

### 18.2.1.17. `.status.loadBalancer.ingress`

#### Description

`ingress` is a list containing ingress points for the load-balancer.

#### Type

**array**

### 18.2.1.18. `.status.loadBalancer.ingress[]`

#### Description

`IngressLoadBalancerIngress` represents the status of a load-balancer ingress point.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>hostname</b>	<b>string</b>	hostname is set for load-balancer ingress points that are DNS based.
<b>ip</b>	<b>string</b>	ip is set for load-balancer ingress points that are IP based.
<b>ports</b>	<b>array</b>	ports provides information about the ports exposed by this LoadBalancer.
<b>ports[]</b>	<b>object</b>	IngressPortStatus represents the error condition of a service port

#### 18.2.1.19. `.status.loadBalancer.ingress[].ports`

##### Description

ports provides information about the ports exposed by this LoadBalancer.

##### Type

**array**

#### 18.2.1.20. `.status.loadBalancer.ingress[].ports[]`

##### Description

IngressPortStatus represents the error condition of a service port

##### Type

**object**

##### Required

- **port**
- **protocol**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>error</b>	<b>string</b>	error is to record the problem with the service port The format of the error shall comply with the following rules: - built-in error values shall be specified in this file and those shall use CamelCase names - cloud provider specific error values must have names that comply with the format foo.example.com/CamelCase.
<b>port</b>	<b>integer</b>	port is the port number of the ingress port.
<b>protocol</b>	<b>string</b>	protocol is the protocol of the ingress port. The supported values are: "TCP", "UDP", "SCTP"  Possible enum values: - " <b>SCTP</b> " is the SCTP protocol. - " <b>TCP</b> " is the TCP protocol. - " <b>UDP</b> " is the UDP protocol.

### 18.2.2. API endpoints

The following API endpoints are available:

- **/apis/networking.k8s.io/v1/ingresses**
  - **GET**: list or watch objects of kind Ingress
- **/apis/networking.k8s.io/v1/watch/ingresses**
  - **GET**: watch individual changes to a list of Ingress. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses**
  - **DELETE**: delete collection of Ingress
  - **GET**: list or watch objects of kind Ingress
  - **POST**: create an Ingress
- **/apis/networking.k8s.io/v1/watch/namespaces/{namespace}/ingresses**
  - **GET**: watch individual changes to a list of Ingress. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}**
  - **DELETE**: delete an Ingress

- **GET**: read the specified Ingress
- **PATCH**: partially update the specified Ingress
- **PUT**: replace the specified Ingress
- **/apis/networking.k8s.io/v1/watch/namespaces/{namespace}/ingresses/{name}**
  - **GET**: watch changes to an object of kind Ingress. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}/status**
  - **GET**: read status of the specified Ingress
  - **PATCH**: partially update status of the specified Ingress
  - **PUT**: replace status of the specified Ingress

### 18.2.2.1. /apis/networking.k8s.io/v1/ingresses

Table 18.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Ingress

**Table 18.2. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>IngressList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 18.2.2.2. /apis/networking.k8s.io/v1/watch/ingresses

Table 18.3. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Ingress. deprecated: use the 'watch' parameter with a list operation instead.

**Table 18.4. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 18.2.2.3. /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses

Table 18.5. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.6. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Ingress

Table 18.7. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 18.8. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 18.9. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Ingress

Table 18.10. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 18.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">IngressList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create an Ingress

**Table 18.12. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 18.13. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Ingress</b> schema	

**Table 18.14. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Ingress</b> schema

HTTP code	Response body
201 - Created	<a href="#">Ingress</a> schema
202 - Accepted	<a href="#">Ingress</a> schema
401 - Unauthorized	Empty

#### 18.2.2.4. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/ingresses

Table 18.15. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.16. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Ingress. deprecated: use the 'watch' parameter with a list operation instead.

Table 18.17. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 18.2.2.5. /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}

Table 18.18. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Ingress
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.19. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete an Ingress

Table 18.20. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 18.21. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 18.22. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Ingress

Table 18.23. HTTP responses

HTTP code	Response body
200 - OK	<b>Ingress</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Ingress

**Table 18.24. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 18.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 18.26. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Ingress</b> schema
201 - Created	<b>Ingress</b> schema
401 - Unauthorized	Empty

## HTTP method

### PUT

#### Description

replace the specified Ingress

Table 18.27. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 18.28. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Ingress</b> schema	

Table 18.29. HTTP responses

HTTP code	Response body
200 - OK	<b>Ingress</b> schema
201 - Created	<b>Ingress</b> schema
401 - Unauthorized	Empty

### 18.2.2.6. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/ingresses/{name}

Table 18.30. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Ingress
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.31. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Ingress. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 18.32. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 18.2.2.7. /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}/status

Table 18.33. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Ingress
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.34. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified Ingress

Table 18.35. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">Ingress</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified Ingress

Table 18.36. Query parameters



Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 18.37. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 18.38. HTTP responses

HTTP code	Response body
200 - OK	<b>Ingress</b> schema
201 - Created	<b>Ingress</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Ingress

**Table 18.39. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 18.40. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Ingress</b> schema	

Table 18.41. HTTP responses

HTTP code	Response body
200 - OK	<b>Ingress</b> schema
201 - Created	<b>Ingress</b> schema
401 - Unauthorized	Empty

## 18.3. INGRESSCLASS [NETWORKING.K8S.IO/V1]

### Description

IngressClass represents the class of the Ingress, referenced by the Ingress Spec. The **ingressclass.kubernetes.io/is-default-class** annotation can be used to indicate that an IngressClass should be considered default. When a single IngressClass resource has this annotation set to true, new Ingress resources without a class specified will be assigned this default class.

### Type

**object**

### 18.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	IngressClassSpec provides information about the class of an Ingress.

### 18.3.1.1. .spec

#### Description

IngressClassSpec provides information about the class of an Ingress.

#### Type

**object**

Property	Type	Description
<b>controller</b>	<b>string</b>	controller refers to the name of the controller that should handle this class. This allows for different "flavors" that are controlled by the same controller. For example, you may have different parameters for the same implementing controller. This should be specified as a domain-prefixed path no more than 250 characters in length, e.g. "acme.io/ingress-controller". This field is immutable.

Property	Type	Description
<b>parameters</b>	<b>object</b>	IngressClassParametersReference identifies an API object. This can be used to specify a cluster or namespace-scoped resource.

### 18.3.1.2. .spec.parameters

#### Description

IngressClassParametersReference identifies an API object. This can be used to specify a cluster or namespace-scoped resource.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	apiGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
<b>kind</b>	<b>string</b>	kind is the type of resource being referenced.
<b>name</b>	<b>string</b>	name is the name of resource being referenced.
<b>namespace</b>	<b>string</b>	namespace is the namespace of the resource being referenced. This field is required when scope is set to "Namespace" and must be unset when scope is set to "Cluster".
<b>scope</b>	<b>string</b>	scope represents if this refers to a cluster or namespace scoped resource. This may be set to "Cluster" (default) or "Namespace".

## 18.3.2. API endpoints

The following API endpoints are available:

- **/apis/networking.k8s.io/v1/ingressclasses**
  - **DELETE**: delete collection of IngressClass
  - **GET**: list or watch objects of kind IngressClass
  - **POST**: create an IngressClass
- **/apis/networking.k8s.io/v1/watch/ingressclasses**
  - **GET**: watch individual changes to a list of IngressClass. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/networking.k8s.io/v1/ingressclasses/{name}**
  - **DELETE**: delete an IngressClass
  - **GET**: read the specified IngressClass
  - **PATCH**: partially update the specified IngressClass
  - **PUT**: replace the specified IngressClass
- **/apis/networking.k8s.io/v1/watch/ingressclasses/{name}**
  - **GET**: watch changes to an object of kind IngressClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 18.3.2.1. /apis/networking.k8s.io/v1/ingressclasses

Table 18.42. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**DELETE**

Description

delete collection of IngressClass

Table 18.43. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 18.44. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 18.45. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind IngressClass

Table 18.46. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 18.47. HTTP responses

HTTP code	Response body
200 - OK	<b>IngressClassList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create an IngressClass

**Table 18.48. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 18.49. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>IngressClass</b> schema	

**Table 18.50. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>IngressClass</b> schema

HTTP code	Response body
201 - Created	<b>IngressClass</b> schema
202 - Accepted	<b>IngressClass</b> schema
401 - Unauthorized	Empty

### 18.3.2.2. /apis/networking.k8s.io/v1/watch/ingressclasses

Table 18.51. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
GET

**Description**

watch individual changes to a list of IngressClass. deprecated: use the 'watch' parameter with a list operation instead.

**Table 18.52. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**18.3.2.3. /apis/networking.k8s.io/v1/ingressclasses/{name}****Table 18.53. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the IngressClass

**Table 18.54. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete an IngressClass

**Table 18.55. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 18.56. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 18.57. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified IngressClass

Table 18.58. HTTP responses

HTTP code	Response body
200 - OK	<b>IngressClass</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified IngressClass

**Table 18.59. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 18.60. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 18.61. HTTP responses

HTTP code	Response body
200 - OK	<b>IngressClass</b> schema
201 - Created	<b>IngressClass</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified IngressClass

Table 18.62. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 18.63. Body parameters

Parameter	Type	Description
<b>body</b>	<b>IngressClass</b> schema	

Table 18.64. HTTP responses

HTTP code	Response body
200 - OK	<b>IngressClass</b> schema
201 - Created	<b>IngressClass</b> schema
401 - Unauthorized	Empty

#### 18.3.2.4. /apis/networking.k8s.io/v1/watch/ingressclasses/{name}

Table 18.65. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the IngressClass

Table 18.66. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind IngressClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 18.67. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 18.4. NETWORKPOLICY [NETWORKING.K8S.IO/V1]

### Description

NetworkPolicy describes what network traffic is allowed for a set of Pods

### Type

**object**

### 18.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	NetworkPolicySpec provides the specification of a NetworkPolicy

Property	Type	Description
<b>status</b>	<b>object</b>	NetworkPolicyStatus describes the current state of the NetworkPolicy.

### 18.4.1.1. .spec

#### Description

NetworkPolicySpec provides the specification of a NetworkPolicy

#### Type

**object**

#### Required

- **podSelector**

Property	Type	Description
<b>egress</b>	<b>array</b>	egress is a list of egress rules to be applied to the selected pods. Outgoing traffic is allowed if there are no NetworkPolicies selecting the pod (and cluster policy otherwise allows the traffic), OR if the traffic matches at least one egress rule across all of the NetworkPolicy objects whose podSelector matches the pod. If this field is empty then this NetworkPolicy limits all outgoing traffic (and serves solely to ensure that the pods it selects are isolated by default). This field is beta-level in 1.8
<b>egress[]</b>	<b>object</b>	NetworkPolicyEgressRule describes a particular set of traffic that is allowed out of pods matched by a NetworkPolicySpec's podSelector. The traffic must match both ports and to. This type is beta-level in 1.8

Property	Type	Description
<b>ingress</b>	<b>array</b>	ingress is a list of ingress rules to be applied to the selected pods. Traffic is allowed to a pod if there are no NetworkPolicies selecting the pod (and cluster policy otherwise allows the traffic), OR if the traffic source is the pod's local node, OR if the traffic matches at least one ingress rule across all of the NetworkPolicy objects whose podSelector matches the pod. If this field is empty then this NetworkPolicy does not allow any traffic (and serves solely to ensure that the pods it selects are isolated by default)
<b>ingress[]</b>	<b>object</b>	NetworkPolicyIngressRule describes a particular set of traffic that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The traffic must match both ports and from.
<b>podSelector</b>	<b>LabelSelector</b>	podSelector selects the pods to which this NetworkPolicy object applies. The array of ingress rules is applied to any pods selected by this field. Multiple network policies can select the same set of pods. In this case, the ingress rules for each are combined additively. This field is NOT optional and follows standard label selector semantics. An empty podSelector matches all pods in this namespace.

Property	Type	Description
<b>policyTypes</b>	<b>array (string)</b>	policyTypes is a list of rule types that the NetworkPolicy relates to. Valid options are ["Ingress"], ["Egress"], or ["Ingress", "Egress"]. If this field is not specified, it will default based on the existence of ingress or egress rules; policies that contain an egress section are assumed to affect egress, and all policies (whether or not they contain an ingress section) are assumed to affect ingress. If you want to write an egress-only policy, you must explicitly specify policyTypes [ "Egress" ]. Likewise, if you want to write a policy that specifies that no egress is allowed, you must specify a policyTypes value that include "Egress" (since such a policy would not include an egress section and would otherwise default to just [ "Ingress" ]). This field is beta-level in 1.8

### 18.4.1.2. .spec.egress

#### Description

egress is a list of egress rules to be applied to the selected pods. Outgoing traffic is allowed if there are no NetworkPolicies selecting the pod (and cluster policy otherwise allows the traffic), OR if the traffic matches at least one egress rule across all of the NetworkPolicy objects whose podSelector matches the pod. If this field is empty then this NetworkPolicy limits all outgoing traffic (and serves solely to ensure that the pods it selects are isolated by default). This field is beta-level in 1.8

#### Type

**array**

### 18.4.1.3. .spec.egress[]

#### Description

NetworkPolicyEgressRule describes a particular set of traffic that is allowed out of pods matched by a NetworkPolicySpec's podSelector. The traffic must match both ports and to. This type is beta-level in 1.8

#### Type

**object**

Property	Type	Description
<b>ports</b>	<b>array</b>	ports is a list of destination ports for outgoing traffic. Each item in this list is combined using a logical OR. If this field is empty or missing, this rule matches all ports (traffic not restricted by port). If this field is present and contains at least one item, then this rule allows traffic only if the traffic matches at least one port in the list.
<b>ports[]</b>	<b>object</b>	NetworkPolicyPort describes a port to allow traffic on
<b>to</b>	<b>array</b>	to is a list of destinations for outgoing traffic of pods selected for this rule. Items in this list are combined using a logical OR operation. If this field is empty or missing, this rule matches all destinations (traffic not restricted by destination). If this field is present and contains at least one item, this rule allows traffic only if the traffic matches at least one item in the to list.
<b>to[]</b>	<b>object</b>	NetworkPolicyPeer describes a peer to allow traffic to/from. Only certain combinations of fields are allowed

#### 18.4.1.4. .spec.egress[].ports

##### Description

ports is a list of destination ports for outgoing traffic. Each item in this list is combined using a logical OR. If this field is empty or missing, this rule matches all ports (traffic not restricted by port). If this field is present and contains at least one item, then this rule allows traffic only if the traffic matches at least one port in the list.

##### Type

**array**

#### 18.4.1.5. .spec.egress[].ports[]

##### Description

NetworkPolicyPort describes a port to allow traffic on

##### Type

**object**

Property	Type	Description
<b>endPort</b>	<b>integer</b>	endPort indicates that the range of ports from port to endPort if set, inclusive, should be allowed by the policy. This field cannot be defined if the port field is not defined or if the port field is defined as a named (string) port. The endPort must be equal or greater than port.
<b>port</b>	<b>IntOrString</b>	port represents the port on the given protocol. This can either be a numerical or named port on a pod. If this field is not provided, this matches all port names and numbers. If present, only traffic on the specified protocol AND port will be matched.
<b>protocol</b>	<b>string</b>	protocol represents the protocol (TCP, UDP, or SCTP) which traffic must match. If not specified, this field defaults to TCP.  Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.

**18.4.1.6. .spec.egress[].to****Description**

to is a list of destinations for outgoing traffic of pods selected for this rule. Items in this list are combined using a logical OR operation. If this field is empty or missing, this rule matches all destinations (traffic not restricted by destination). If this field is present and contains at least one item, this rule allows traffic only if the traffic matches at least one item in the to list.

**Type**

**array**

**18.4.1.7. .spec.egress[].to[]****Description**

NetworkPolicyPeer describes a peer to allow traffic to/from. Only certain combinations of fields are allowed

**Type**

**object**

Property	Type	Description
<b>ipBlock</b>	<b>object</b>	IPBlock describes a particular CIDR (Ex. "192.168.1.0/24","2001:db8::/64") that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The except entry describes CIDRs that should not be included within this rule.
<b>namespaceSelector</b>	<b>LabelSelector</b>	namespaceSelector selects namespaces using cluster-scoped labels. This field follows standard label selector semantics; if present but empty, it selects all namespaces.  If podSelector is also set, then the NetworkPolicyPeer as a whole selects the pods matching podSelector in the namespaces selected by namespaceSelector. Otherwise it selects all pods in the namespaces selected by namespaceSelector.
<b>podSelector</b>	<b>LabelSelector</b>	podSelector is a label selector which selects pods. This field follows standard label selector semantics; if present but empty, it selects all pods.  If namespaceSelector is also set, then the NetworkPolicyPeer as a whole selects the pods matching podSelector in the Namespaces selected by NamespaceSelector. Otherwise it selects the pods matching podSelector in the policy's own namespace.

#### 18.4.1.8. .spec.egress[].to[].ipBlock

##### Description

IPBlock describes a particular CIDR (Ex. "192.168.1.0/24","2001:db8::/64") that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The except entry describes CIDRs that should not be included within this rule.

##### Type

**object**

**Required**



- **cidr**

Property	Type	Description
<b>cidr</b>	<b>string</b>	cidr is a string representing the IPBlock Valid examples are "192.168.1.0/24" or "2001:db8::/64"
<b>except</b>	<b>array (string)</b>	except is a slice of CIDRs that should not be included within an IPBlock Valid examples are "192.168.1.0/24" or "2001:db8::/64" Except values will be rejected if they are outside the cidr range

#### 18.4.1.9. .spec.ingress

##### Description

ingress is a list of ingress rules to be applied to the selected pods. Traffic is allowed to a pod if there are no NetworkPolicies selecting the pod (and cluster policy otherwise allows the traffic), OR if the traffic source is the pod's local node, OR if the traffic matches at least one ingress rule across all of the NetworkPolicy objects whose podSelector matches the pod. If this field is empty then this NetworkPolicy does not allow any traffic (and serves solely to ensure that the pods it selects are isolated by default)

##### Type

**array**

#### 18.4.1.10. .spec.ingress[]

##### Description

NetworkPolicyIngressRule describes a particular set of traffic that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The traffic must match both ports and from.

##### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>from</b>	<b>array</b>	from is a list of sources which should be able to access the pods selected for this rule. Items in this list are combined using a logical OR operation. If this field is empty or missing, this rule matches all sources (traffic not restricted by source). If this field is present and contains at least one item, this rule allows traffic only if the traffic matches at least one item in the from list.
<b>from[]</b>	<b>object</b>	NetworkPolicyPeer describes a peer to allow traffic to/from. Only certain combinations of fields are allowed
<b>ports</b>	<b>array</b>	ports is a list of ports which should be made accessible on the pods selected for this rule. Each item in this list is combined using a logical OR. If this field is empty or missing, this rule matches all ports (traffic not restricted by port). If this field is present and contains at least one item, then this rule allows traffic only if the traffic matches at least one port in the list.
<b>ports[]</b>	<b>object</b>	NetworkPolicyPort describes a port to allow traffic on

#### 18.4.1.11. .spec.ingress[].from

##### Description

from is a list of sources which should be able to access the pods selected for this rule. Items in this list are combined using a logical OR operation. If this field is empty or missing, this rule matches all sources (traffic not restricted by source). If this field is present and contains at least one item, this rule allows traffic only if the traffic matches at least one item in the from list.

##### Type

**array**

#### 18.4.1.12. .spec.ingress[].from[]

##### Description

NetworkPolicyPeer describes a peer to allow traffic to/from. Only certain combinations of fields are allowed

## Type

**object**

Property	Type	Description
<b>ipBlock</b>	<b>object</b>	IPBlock describes a particular CIDR (Ex. "192.168.1.0/24", "2001:db8::/64") that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The except entry describes CIDRs that should not be included within this rule.
<b>namespaceSelector</b>	<a href="#">LabelSelector</a>	namespaceSelector selects namespaces using cluster-scoped labels. This field follows standard label selector semantics; if present but empty, it selects all namespaces.  If podSelector is also set, then the NetworkPolicyPeer as a whole selects the pods matching podSelector in the namespaces selected by namespaceSelector. Otherwise it selects all pods in the namespaces selected by namespaceSelector.
<b>podSelector</b>	<a href="#">LabelSelector</a>	podSelector is a label selector which selects pods. This field follows standard label selector semantics; if present but empty, it selects all pods.  If namespaceSelector is also set, then the NetworkPolicyPeer as a whole selects the pods matching podSelector in the Namespaces selected by NamespaceSelector. Otherwise it selects the pods matching podSelector in the policy's own namespace.

### 18.4.1.13. .spec.ingress[].from[].ipBlock

#### Description

IPBlock describes a particular CIDR (Ex. "192.168.1.0/24", "2001:db8::/64") that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The except entry describes CIDRs that should not be included within this rule.

IPBlock describes a particular CIDR (Ex. "192.168.1.0/24", "2001:db8::/64") that is allowed to the pods matched by a NetworkPolicySpec's podSelector. The except entry describes CIDRs that should not be included within this rule.

**Type****object****Required**

- **cidr**

Property	Type	Description
<b>cidr</b>	<b>string</b>	cidr is a string representing the IPBlock Valid examples are "192.168.1.0/24" or "2001:db8::/64"
<b>except</b>	<b>array (string)</b>	except is a slice of CIDRs that should not be included within an IPBlock Valid examples are "192.168.1.0/24" or "2001:db8::/64" Except values will be rejected if they are outside the cidr range

**18.4.1.14. .spec.ingress[].ports****Description**

ports is a list of ports which should be made accessible on the pods selected for this rule. Each item in this list is combined using a logical OR. If this field is empty or missing, this rule matches all ports (traffic not restricted by port). If this field is present and contains at least one item, then this rule allows traffic only if the traffic matches at least one port in the list.

**Type****array****18.4.1.15. .spec.ingress[].ports[]****Description**

NetworkPolicyPort describes a port to allow traffic on

**Type****object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>endPort</b>	<b>integer</b>	endPort indicates that the range of ports from port to endPort if set, inclusive, should be allowed by the policy. This field cannot be defined if the port field is not defined or if the port field is defined as a named (string) port. The endPort must be equal or greater than port.
<b>port</b>	<b>IntOrString</b>	port represents the port on the given protocol. This can either be a numerical or named port on a pod. If this field is not provided, this matches all port names and numbers. If present, only traffic on the specified protocol AND port will be matched.
<b>protocol</b>	<b>string</b>	protocol represents the protocol (TCP, UDP, or SCTP) which traffic must match. If not specified, this field defaults to TCP.  Possible enum values: - <b>"SCTP"</b> is the SCTP protocol. - <b>"TCP"</b> is the TCP protocol. - <b>"UDP"</b> is the UDP protocol.

#### 18.4.1.16. .status

##### Description

NetworkPolicyStatus describes the current state of the NetworkPolicy.

##### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array (Condition)</b>	conditions holds an array of metav1.Condition that describe the state of the NetworkPolicy. Current service state

#### 18.4.2. API endpoints

The following API endpoints are available:

- **/apis/networking.k8s.io/v1/networkpolicies**
  - **GET:** list or watch objects of kind NetworkPolicy
- **/apis/networking.k8s.io/v1/watch/networkpolicies**
  - **GET:** watch individual changes to a list of NetworkPolicy. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies**
  - **DELETE:** delete collection of NetworkPolicy
  - **GET:** list or watch objects of kind NetworkPolicy
  - **POST:** create a NetworkPolicy
- **/apis/networking.k8s.io/v1/watch/namespaces/{namespace}/networkpolicies**
  - **GET:** watch individual changes to a list of NetworkPolicy. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}**
  - **DELETE:** delete a NetworkPolicy
  - **GET:** read the specified NetworkPolicy
  - **PATCH:** partially update the specified NetworkPolicy
  - **PUT:** replace the specified NetworkPolicy
- **/apis/networking.k8s.io/v1/watch/namespaces/{namespace}/networkpolicies/{name}**
  - **GET:** watch changes to an object of kind NetworkPolicy. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}/status**
  - **GET:** read status of the specified NetworkPolicy
  - **PATCH:** partially update status of the specified NetworkPolicy
  - **PUT:** replace status of the specified NetworkPolicy

#### 18.4.2.1. /apis/networking.k8s.io/v1/networkpolicies

Table 18.68. Global query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind NetworkPolicy

Table 18.69. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">NetworkPolicyList</a> schema
401 - Unauthorized	Empty

#### 18.4.2.2. /apis/networking.k8s.io/v1/watch/networkpolicies

Table 18.70. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of NetworkPolicy. deprecated: use the 'watch' parameter with a list operation instead.

**Table 18.71. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 18.4.2.3. /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies

Table 18.72. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.73. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of NetworkPolicy

Table 18.74. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>



Table 18.75. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 18.76. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind NetworkPolicy

Table 18.77. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 18.78. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">NetworkPolicyList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a NetworkPolicy

**Table 18.79. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 18.80. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>NetworkPolicy</b> schema	

**Table 18.81. HTTP responses**

HTTP code	Response body
200 - OK	<b>NetworkPolicy</b> schema

HTTP code	Response body
201 - Created	<a href="#">NetworkPolicy</a> schema
202 - Accepted	<a href="#">NetworkPolicy</a> schema
401 - Unauthorized	Empty

#### 18.4.2.4. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/networkpolicies

Table 18.82. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.83. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of NetworkPolicy. deprecated: use the 'watch' parameter with a list operation instead.

Table 18.84. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 18.4.2.5. /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}

Table 18.85. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the NetworkPolicy
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.86. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete a NetworkPolicy

Table 18.87. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 18.88. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 18.89. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified NetworkPolicy

Table 18.90. HTTP responses

HTTP code	Response body
200 - OK	<b>NetworkPolicy</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified NetworkPolicy

**Table 18.91. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 18.92. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 18.93. HTTP responses

HTTP code	Response body
200 - OK	<b>NetworkPolicy</b> schema
201 - Created	<b>NetworkPolicy</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified NetworkPolicy

Table 18.94. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 18.95. Body parameters

Parameter	Type	Description
<b>body</b>	<b>NetworkPolicy</b> schema	

Table 18.96. HTTP responses

HTTP code	Response body
200 - OK	<b>NetworkPolicy</b> schema
201 - Created	<b>NetworkPolicy</b> schema
401 - Unauthorized	Empty

#### 18.4.2.6. /apis/networking.k8s.io/v1/watch/namespaces/{namespace}/networkpolicies/{name}

Table 18.97. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the NetworkPolicy
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 18.98. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind NetworkPolicy. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 18.99. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

#### 18.4.2.7. /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}/status

Table 18.100. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the NetworkPolicy
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 18.101. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified NetworkPolicy

Table 18.102. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">NetworkPolicy</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified NetworkPolicy

Table 18.103. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 18.104. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 18.105. HTTP responses

HTTP code	Response body
200 - OK	<b>NetworkPolicy</b> schema
201 - Created	<b>NetworkPolicy</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified NetworkPolicy

**Table 18.106. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 18.107. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>NetworkPolicy</b> schema	

Table 18.108. HTTP responses

HTTP code	Response body
200 - OK	<b>NetworkPolicy</b> schema
201 - Created	<b>NetworkPolicy</b> schema
401 - Unauthorized	Empty

## CHAPTER 19. NODE APIS

### 19.1. NODE APIS

#### 19.1.1. RuntimeClass [node.k8s.io/v1]

##### Description

RuntimeClass defines a class of container runtime supported in the cluster. The RuntimeClass is used to determine which container runtime is used to run all containers in a pod. RuntimeClasses are manually defined by a user or cluster provisioner, and referenced in the PodSpec. The Kubelet is responsible for resolving the RuntimeClassName reference before running the pod. For more details, see <https://kubernetes.io/docs/concepts/containers/runtime-class/>

##### Type

**object**

### 19.2. RUNTIMECLASS [NODE.K8S.IO/V1]

##### Description

RuntimeClass defines a class of container runtime supported in the cluster. The RuntimeClass is used to determine which container runtime is used to run all containers in a pod. RuntimeClasses are manually defined by a user or cluster provisioner, and referenced in the PodSpec. The Kubelet is responsible for resolving the RuntimeClassName reference before running the pod. For more details, see <https://kubernetes.io/docs/concepts/containers/runtime-class/>

##### Type

**object**

##### Required

- **handler**

#### 19.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>handler</b>	<b>string</b>	handler specifies the underlying runtime and configuration that the CRI implementation will use to handle pods of this class. The possible values are specific to the node & CRI configuration. It is assumed that all handlers are available on every node, and handlers of the same name are equivalent on every node. For example, a handler called "runc" might specify that the runc OCI runtime (using native Linux containers) will be used to run the containers in a pod. The Handler must be lowercase, conform to the DNS Label (RFC 1123) requirements, and is immutable.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>overhead</b>	<b>object</b>	Overhead structure represents the resource overhead associated with running a pod.
<b>scheduling</b>	<b>object</b>	Scheduling specifies the scheduling constraints for nodes supporting a RuntimeClass.

### 19.2.1.1. .overhead

#### Description

Overhead structure represents the resource overhead associated with running a pod.

#### Type

**object**

Property	Type	Description
<b>podFixed</b>	<b>object (Quantity)</b>	podFixed represents the fixed resource overhead associated with running a pod.

**19.2.1.2. .scheduling****Description**

Scheduling specifies the scheduling constraints for nodes supporting a RuntimeClass.

**Type****object**

Property	Type	Description
<b>nodeSelector</b>	<b>object (string)</b>	nodeSelector lists labels that must be present on nodes that support this RuntimeClass. Pods using this RuntimeClass can only be scheduled to a node matched by this selector. The RuntimeClass nodeSelector is merged with a pod's existing nodeSelector. Any conflicts will cause the pod to be rejected in admission.
<b>tolerations</b>	<b>array (Toleration)</b>	tolerations are appended (excluding duplicates) to pods running with this RuntimeClass during admission, effectively unioning the set of nodes tolerated by the pod and the RuntimeClass.

**19.2.2. API endpoints**

The following API endpoints are available:

- **/apis/node.k8s.io/v1/runtimeclasses**
  - **DELETE**: delete collection of RuntimeClass
  - **GET**: list or watch objects of kind RuntimeClass
  - **POST**: create a RuntimeClass
- **/apis/node.k8s.io/v1/watch/runtimeclasses**



- **GET**: watch individual changes to a list of RuntimeClass. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/node.k8s.io/v1/runtimeclasses/{name}**
  - **DELETE**: delete a RuntimeClass
  - **GET**: read the specified RuntimeClass
  - **PATCH**: partially update the specified RuntimeClass
  - **PUT**: replace the specified RuntimeClass
- **/apis/node.k8s.io/v1/watch/runtimeclasses/{name}**
  - **GET**: watch changes to an object of kind RuntimeClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 19.2.2.1. /apis/node.k8s.io/v1/runtimeclasses

Table 19.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of RuntimeClass

Table 19.2. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 19.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 19.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind RuntimeClass

Table 19.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 19.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">RuntimeClassList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**



**Description**

create a `RuntimeClass`

**Table 19.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized <code>dryRun</code> directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	<code>fieldManager</code> is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	<code>fieldValidation</code> instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a <code>BadRequest</code> error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 19.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>RuntimeClass</b> schema	

**Table 19.9. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>RuntimeClass</b> schema

HTTP code	Response body
201 - Created	<a href="#">RuntimeClass</a> schema
202 - Accepted	<a href="#">RuntimeClass</a> schema
401 - Unauthorized	Empty

### 19.2.2.2. /apis/node.k8s.io/v1/watch/runtimeclasses

Table 19.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
**GET**

**Description**

watch individual changes to a list of RuntimeClass. deprecated: use the 'watch' parameter with a list operation instead.

**Table 19.11. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**19.2.2.3. /apis/node.k8s.io/v1/runtimeclasses/{name}****Table 19.12. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the RuntimeClass

**Table 19.13. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a RuntimeClass

**Table 19.14. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 19.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 19.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified RuntimeClass

Table 19.17. HTTP responses

HTTP code	Response body
200 - OK	<b>RuntimeClass</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified RuntimeClass

**Table 19.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 19.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 19.20. HTTP responses

HTTP code	Response body
200 - OK	<b>RuntimeClass</b> schema
201 - Created	<b>RuntimeClass</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified RuntimeClass

Table 19.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 19.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>RuntimeClass</b> schema	

Table 19.23. HTTP responses

HTTP code	Response body
200 - OK	<b>RuntimeClass</b> schema
201 - Created	<b>RuntimeClass</b> schema
401 - Unauthorized	Empty

#### 19.2.2.4. /apis/node.k8s.io/v1/watch/runtimeclasses/{name}

Table 19.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the RuntimeClass

Table 19.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind `RuntimeClass`. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 19.26. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## CHAPTER 20. POLICY APIS

### 20.1. POLICY APIS

#### 20.1.1. Eviction [policy/v1]

##### Description

Eviction evicts a pod from its node subject to certain policies and safety constraints. This is a subresource of Pod. A request to cause such an eviction is created by POSTing to `.../pods/<pod name>/evictions`.

##### Type

**object**

#### 20.1.2. PodDisruptionBudget [policy/v1]

##### Description

PodDisruptionBudget is an object to define the max disruption that can be caused to a collection of pods

##### Type

**object**

### 20.2. EVICTION [POLICY/V1]

##### Description

Eviction evicts a pod from its node subject to certain policies and safety constraints. This is a subresource of Pod. A request to cause such an eviction is created by POSTing to `.../pods/<pod name>/evictions`.

##### Type

**object**

#### 20.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>deleteOptions</b>	<b>DeleteOptions</b>	DeleteOptions may be provided

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	ObjectMeta describes the pod that is being evicted.

## 20.2.2. API endpoints

The following API endpoints are available:

- **/api/v1/namespaces/{namespace}/pods/{name}/eviction**
  - **POST**: create eviction of a Pod

### 20.2.2.1. /api/v1/namespaces/{namespace}/pods/{name}/eviction

Table 20.1. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Eviction
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 20.2. Global query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****POST****Description**

create eviction of a Pod

**Table 20.3. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Eviction</b> schema	

**Table 20.4. HTTP responses**

HTTP code	Response body
200 - OK	<b>Eviction</b> schema
201 - Created	<b>Eviction</b> schema



HTTP code	Reponse body
202 - Accepted	<b>Eviction</b> schema
401 - Unauthorized	Empty

## 20.3. PODDISRUPTIONBUDGET [POLICY/V1]

### Description

PodDisruptionBudget is an object to define the max disruption that can be caused to a collection of pods

### Type

**object**

### 20.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>spec</b>	<b>object</b>	PodDisruptionBudgetSpec is a description of a PodDisruptionBudget.
<b>status</b>	<b>object</b>	PodDisruptionBudgetStatus represents information about the status of a PodDisruptionBudget. Status may trail the actual state of a system.

### 20.3.1.1. .spec

#### Description

PodDisruptionBudgetSpec is a description of a PodDisruptionBudget.

#### Type

**object**

Property	Type	Description
<b>maxUnavailable</b>	<b>IntOrString</b>	An eviction is allowed if at most "maxUnavailable" pods selected by "selector" are unavailable after the eviction, i.e. even in absence of the evicted pod. For example, one can prevent all voluntary evictions by specifying 0. This is a mutually exclusive setting with "minAvailable".
<b>minAvailable</b>	<b>IntOrString</b>	An eviction is allowed if at least "minAvailable" pods selected by "selector" will still be available after the eviction, i.e. even in the absence of the evicted pod. So for example you can prevent all voluntary evictions by specifying "100%".
<b>selector</b>	<b>LabelSelector</b>	Label query over pods whose evictions are managed by the disruption budget. A null selector will match no pods, while an empty ({} ) selector will select all pods within the namespace.
<b>unhealthyPodEvictionPolicy</b>	<b>string</b>	UnhealthyPodEvictionPolicy

Property	Type	Description
		<p>defines the criteria for when unhealthy pods should be considered for eviction. Current implementation considers healthy pods, as pods that have <code>status.conditions</code> item with <code>type="Ready",status="True"</code>.</p> <p>Valid policies are <code>IfHealthyBudget</code> and <code>AlwaysAllow</code>. If no policy is specified, the default behavior will be used, which corresponds to the <code>IfHealthyBudget</code> policy.</p> <p><code>IfHealthyBudget</code> policy means that running pods (<code>status.phase="Running"</code>), but not yet healthy can be evicted only if the guarded application is not disrupted (<code>status.currentHealthy</code> is at least equal to <code>status.desiredHealthy</code>). Healthy pods will be subject to the PDB for eviction.</p> <p><code>AlwaysAllow</code> policy means that all running pods (<code>status.phase="Running"</code>), but not yet healthy are considered disrupted and can be evicted regardless of whether the criteria in a PDB is met. This means perspective running pods of a disrupted application might not get a chance to become healthy. Healthy pods will be subject to the PDB for eviction.</p> <p>Additional policies may be added in the future. Clients making eviction decisions should disallow eviction of unhealthy pods if they encounter an unrecognized policy in this field.</p> <p>This field is beta-level. The eviction API uses this field when the feature gate <code>PDBUnhealthyPodEvictionPolicy</code> is enabled (enabled by default).</p> <p>Possible enum values: -  <b>"AlwaysAllow"</b> policy means that all running pods (<code>status.phase="Running"</code>), but not yet healthy are considered disrupted and can be evicted</p>

Property	Type	Description
		<p>regardless of whether the criteria in a PDB is met. This means perspective running pods of a disrupted application might not get a chance to become healthy. Healthy pods will be subject to the PDB for eviction. - <b>"IfHealthyBudget"</b> policy means that running pods (status.phase="Running"), but not yet healthy can be evicted only if the guarded application is not disrupted (status.currentHealthy is at least equal to status.desiredHealthy). Healthy pods will be subject to the PDB for eviction.</p>

### 20.3.1.2. .status

#### Description

PodDisruptionBudgetStatus represents information about the status of a PodDisruptionBudget. Status may trail the actual state of a system.

#### Type

**object**

#### Required

- **disruptionsAllowed**
- **currentHealthy**
- **desiredHealthy**
- **expectedPods**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>conditions</b>	<b>array (Condition)</b>	Conditions contain conditions for PDB. The disruption controller sets the DisruptionAllowed condition. The following are known values for the reason field (additional reasons could be added in the future): - SyncFailed: The controller encountered an error and wasn't able to compute the number of allowed disruptions. Therefore no disruptions are allowed and the status of the condition will be False. - InsufficientPods: The number of pods are either at or below the number required by the PodDisruptionBudget. No disruptions are allowed and the status of the condition will be False. - SufficientPods: There are more pods than required by the PodDisruptionBudget. The condition will be True, and the number of allowed disruptions are provided by the disruptionsAllowed property.
<b>currentHealthy</b>	<b>integer</b>	current number of healthy pods
<b>desiredHealthy</b>	<b>integer</b>	minimum desired number of healthy pods

Property	Type	Description
<b>disruptedPods</b>	<b>object (Time)</b>	DisruptedPods contains information about pods whose eviction was processed by the API server eviction subresource handler but has not yet been observed by the PodDisruptionBudget controller. A pod will be in this map from the time when the API server processed the eviction request to the time when the pod is seen by PDB controller as having been marked for deletion (or after a timeout). The key in the map is the name of the pod and the value is the time when the API server processed the eviction request. If the deletion didn't occur and a pod is still there it will be removed from the list automatically by PodDisruptionBudget controller after some time. If everything goes smooth this map should be empty for the most of the time. Large number of entries in the map may indicate problems with pod deletions.
<b>disruptionsAllowed</b>	<b>integer</b>	Number of pod disruptions that are currently allowed.
<b>expectedPods</b>	<b>integer</b>	total number of pods counted by this disruption budget
<b>observedGeneration</b>	<b>integer</b>	Most recent generation observed when updating this PDB status. DisruptionsAllowed and other status information is valid only if observedGeneration equals to PDB's object generation.

### 20.3.2. API endpoints

The following API endpoints are available:

- **/apis/policy/v1/poddisruptionbudgets**
  - **GET**: list or watch objects of kind PodDisruptionBudget

- **/apis/policy/v1/watch/poddisruptionbudgets**
  - **GET**: watch individual changes to a list of PodDisruptionBudget. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets**
  - **DELETE**: delete collection of PodDisruptionBudget
  - **GET**: list or watch objects of kind PodDisruptionBudget
  - **POST**: create a PodDisruptionBudget
- **/apis/policy/v1/watch/namespaces/{namespace}/poddisruptionbudgets**
  - **GET**: watch individual changes to a list of PodDisruptionBudget. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets/{name}**
  - **DELETE**: delete a PodDisruptionBudget
  - **GET**: read the specified PodDisruptionBudget
  - **PATCH**: partially update the specified PodDisruptionBudget
  - **PUT**: replace the specified PodDisruptionBudget
- **/apis/policy/v1/watch/namespaces/{namespace}/poddisruptionbudgets/{name}**
  - **GET**: watch changes to an object of kind PodDisruptionBudget. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets/{name}/status**
  - **GET**: read status of the specified PodDisruptionBudget
  - **PATCH**: partially update status of the specified PodDisruptionBudget
  - **PUT**: replace status of the specified PodDisruptionBudget

### 20.3.2.1. /apis/policy/v1/poddisruptionbudgets

Table 20.5. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind PodDisruptionBudget

Table 20.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PodDisruptionBudgetList</a> schema
401 - Unauthorized	Empty

### 20.3.2.2. /apis/policy/v1/watch/poddisruptionbudgets

Table 20.7. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PodDisruptionBudget. deprecated: use the 'watch' parameter with a list operation instead.

**Table 20.8. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 20.3.2.3. /apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets

Table 20.9. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 20.10. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of PodDisruptionBudget

Table 20.11. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 20.12. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 20.13. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind PodDisruptionBudget

Table 20.14. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 20.15. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PodDisruptionBudgetList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a PodDisruptionBudget

**Table 20.16. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 20.17. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PodDisruptionBudget</b> schema	

**Table 20.18. HTTP responses**

HTTP code	Response body
200 - OK	<b>PodDisruptionBudget</b> schema

HTTP code	Reponse body
201 - Created	<b>PodDisruptionBudget</b> schema
202 - Accepted	<b>PodDisruptionBudget</b> schema
401 - Unauthorized	Empty

#### 20.3.2.4. /apis/policy/v1/watch/namespaces/{namespace}/poddisruptionbudgets

Table 20.19. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 20.20. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of PodDisruptionBudget. deprecated: use the 'watch' parameter with a list operation instead.

Table 20.21. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 20.3.2.5. /apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets/{name}

Table 20.22. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PodDisruptionBudget
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 20.23. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a PodDisruptionBudget

Table 20.24. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 20.25. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 20.26. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified PodDisruptionBudget

Table 20.27. HTTP responses

HTTP code	Response body
200 - OK	<b>PodDisruptionBudget</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified PodDisruptionBudget

**Table 20.28. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 20.29. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 20.30. HTTP responses

HTTP code	Response body
200 - OK	<b>PodDisruptionBudget</b> schema
201 - Created	<b>PodDisruptionBudget</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified PodDisruptionBudget

Table 20.31. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 20.32. Body parameters

Parameter	Type	Description
<b>body</b>	<b>PodDisruptionBudget</b> t schema	

Table 20.33. HTTP responses

HTTP code	Response body
200 - OK	<b>PodDisruptionBudget</b> schema
201 - Created	<b>PodDisruptionBudget</b> schema
401 - Unauthorized	Empty

### 20.3.2.6. /apis/policy/v1/watch/namespaces/{namespace}/poddisruptionbudgets/{name}

Table 20.34. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PodDisruptionBudget
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 20.35. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind PodDisruptionBudget. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 20.36. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 20.3.2.7. /apis/policy/v1/namespaces/{namespace}/poddisruptionbudgets/{name}/status

Table 20.37. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PodDisruptionBudget
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 20.38. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified PodDisruptionBudget

Table 20.39. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">PodDisruptionBudget</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified PodDisruptionBudget

Table 20.40. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 20.41. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 20.42. HTTP responses

HTTP code	Response body
200 - OK	<b>PodDisruptionBudget</b> schema
201 - Created	<b>PodDisruptionBudget</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified PodDisruptionBudget

**Table 20.43. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 20.44. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PodDisruptionBudget</b> t schema	

Table 20.45. HTTP responses

HTTP code	Reponse body
200 - OK	<b>PodDisruptionBudget</b> schema
201 - Created	<b>PodDisruptionBudget</b> schema
401 - Unauthorized	Empty

## CHAPTER 21. RBAC APIS

### 21.1. RBAC APIS

#### 21.1.1. ClusterRole [rbac.authorization.k8s.io/v1]

##### Description

ClusterRole is a cluster level, logical grouping of PolicyRules that can be referenced as a unit by a RoleBinding or ClusterRoleBinding.

##### Type

**object**

#### 21.1.2. ClusterRoleBinding [rbac.authorization.k8s.io/v1]

##### Description

ClusterRoleBinding references a ClusterRole, but not contain it. It can reference a ClusterRole in the global namespace, and adds who information via Subject.

##### Type

**object**

#### 21.1.3. Role [rbac.authorization.k8s.io/v1]

##### Description

Role is a namespaced, logical grouping of PolicyRules that can be referenced as a unit by a RoleBinding.

##### Type

**object**

#### 21.1.4. RoleBinding [rbac.authorization.k8s.io/v1]

##### Description

RoleBinding references a role, but does not contain it. It can reference a Role in the same namespace or a ClusterRole in the global namespace. It adds who information via Subjects and namespace information by which namespace it exists in. RoleBindings in a given namespace only have effect in that namespace.

##### Type

**object**

### 21.2. CLUSTERROLE [RBAC.AUTHORIZATION.K8S.IO/V1]

##### Description

ClusterRole is a cluster level, logical grouping of PolicyRules that can be referenced as a unit by a RoleBinding or ClusterRoleBinding.

##### Type

**object**

#### 21.2.1. Specification

Property	Type	Description
<b>aggregationRule</b>	<b>object</b>	AggregationRule describes how to locate ClusterRoles to aggregate into the ClusterRole
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata.
<b>rules</b>	<b>array</b>	Rules holds all the PolicyRules for this ClusterRole
<b>rules[]</b>	<b>object</b>	PolicyRule holds information that describes a policy rule, but does not contain information about who the rule applies to or which namespace the rule applies to.

### 21.2.1.1. .aggregationRule

#### Description

AggregationRule describes how to locate ClusterRoles to aggregate into the ClusterRole

#### Type

**object**

Property	Type	Description
----------	------	-------------



Property	Type	Description
<b>clusterRoleSelectors</b>	<b>array (LabelSelector)</b>	ClusterRoleSelectors holds a list of selectors which will be used to find ClusterRoles and create the rules. If any of the selectors match, then the ClusterRole's permissions will be added

### 21.2.1.2. .rules

#### Description

Rules holds all the PolicyRules for this ClusterRole

#### Type

**array**

### 21.2.1.3. .rules[]

#### Description

PolicyRule holds information that describes a policy rule, but does not contain information about who the rule applies to or which namespace the rule applies to.

#### Type

**object**

#### Required

- **verbs**

Property	Type	Description
<b>apiGroups</b>	<b>array (string)</b>	APIGroups is the name of the APIGroup that contains the resources. If multiple API groups are specified, any action requested against one of the enumerated resources in any API group will be allowed. "" represents the core API group and "*" represents all API groups.

Property	Type	Description
<b>nonResourceURLs</b>	<b>array (string)</b>	NonResourceURLs is a set of partial urls that a user should have access to. *s are allowed, but only as the full, final step in the path Since non-resource URLs are not namespaced, this field is only applicable for ClusterRoles referenced from a ClusterRoleBinding. Rules can either apply to API resources (such as "pods" or "secrets") or non-resource URL paths (such as "/api"), but not both.
<b>resourceNames</b>	<b>array (string)</b>	ResourceNames is an optional white list of names that the rule applies to. An empty set means that everything is allowed.
<b>resources</b>	<b>array (string)</b>	Resources is a list of resources this rule applies to. '*' represents all resources.
<b>verbs</b>	<b>array (string)</b>	Verbs is a list of Verbs that apply to ALL the ResourceKinds contained in this rule. '*' represents all verbs.

### 21.2.2. API endpoints

The following API endpoints are available:

- **/apis/rbac.authorization.k8s.io/v1/clusterroles**
  - **DELETE**: delete collection of ClusterRole
  - **GET**: list or watch objects of kind ClusterRole
  - **POST**: create a ClusterRole
- **/apis/rbac.authorization.k8s.io/v1/watch/clusterroles**
  - **GET**: watch individual changes to a list of ClusterRole. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/rbac.authorization.k8s.io/v1/clusterroles/{name}**
  - **DELETE**: delete a ClusterRole
  - **GET**: read the specified ClusterRole

- **PATCH**: partially update the specified ClusterRole
- **PUT**: replace the specified ClusterRole
- **/apis/rbac.authorization.k8s.io/v1/watch/clusterroles/{name}**
  - **GET**: watch changes to an object of kind ClusterRole. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 21.2.2.1. /apis/rbac.authorization.k8s.io/v1/clusterroles

Table 21.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of ClusterRole

Table 21.2. Query parameters

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 21.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ClusterRole

**Table 21.5. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 21.6. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRoleList</b> schema
401 - Unauthorized	Empty

## HTTP method

**POST****Description**

create a ClusterRole

**Table 21.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 21.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ClusterRole</b> schema	

**Table 21.9. HTTP responses**

HTTP code	Response body
200 - OK	<b>ClusterRole</b> schema

HTTP code	Response body
201 - Created	<b>ClusterRole</b> schema
202 - Accepted	<b>ClusterRole</b> schema
401 - Unauthorized	Empty

### 21.2.2.2. /apis/rbac.authorization.k8s.io/v1/watch/clusterroles

Table 21.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ClusterRole. deprecated: use the 'watch' parameter with a list operation instead.

Table 21.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 21.2.2.3. /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}

Table 21.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ClusterRole

Table 21.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a ClusterRole

Table 21.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 21.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ClusterRole

Table 21.17. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRole</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ClusterRole

**Table 21.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.



Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 21.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 21.20. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRole</b> schema
201 - Created	<b>ClusterRole</b> schema
401 - Unauthorized	Empty

## HTTP method

### PUT

#### Description

replace the specified ClusterRole

Table 21.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 21.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ClusterRole</b> schema	

Table 21.23. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRole</b> schema
201 - Created	<b>ClusterRole</b> schema
401 - Unauthorized	Empty

#### 21.2.2.4. /apis/rbac.authorization.k8s.io/v1/watch/clusterroles/{name}

Table 21.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ClusterRole

Table 21.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ClusterRole. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 21.26. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 21.3. CLUSTERROLEBINDING [RBAC.AUTHORIZATION.K8S.IO/V1]

### Description

ClusterRoleBinding references a ClusterRole, but not contain it. It can reference a ClusterRole in the global namespace, and adds who information via Subject.

### Type

**object**

### Required

- **roleRef**

### 21.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata.
<b>roleRef</b>	<b>object</b>	RoleRef contains information that points to the role being used

Property	Type	Description
<b>subjects</b>	<b>array</b>	Subjects holds references to the objects the role applies to.
<b>subjects[]</b>	<b>object</b>	Subject contains a reference to the object or user identities a role binding applies to. This can either hold a direct API object reference, or a value for non-objects such as user and group names.

### 21.3.1.1. .roleRef

#### Description

RoleRef contains information that points to the role being used

#### Type

**object**

#### Required

- **apiGroup**
- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

### 21.3.1.2. .subjects

#### Description

Subjects holds references to the objects the role applies to.

#### Type

**array**

### 21.3.1.3. .subjects[]

#### Description

Subject contains a reference to the object or user identities a role binding applies to. This can either hold a direct API object reference, or a value for non-objects such as user and group names.

### Type

**object**

### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
<b>kind</b>	<b>string</b>	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
<b>name</b>	<b>string</b>	Name of the object being referenced.
<b>namespace</b>	<b>string</b>	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

## 21.3.2. API endpoints

The following API endpoints are available:

- **/apis/rbac.authorization.k8s.io/v1/clusterrolebindings**
  - **DELETE**: delete collection of ClusterRoleBinding
  - **GET**: list or watch objects of kind ClusterRoleBinding
  - **POST**: create a ClusterRoleBinding
- **/apis/rbac.authorization.k8s.io/v1/watch/clusterrolebindings**
  - **GET**: watch individual changes to a list of ClusterRoleBinding. deprecated: use the 'watch'



parameter with a list operation instead.

- **/apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}**
  - **DELETE:** delete a ClusterRoleBinding
  - **GET:** read the specified ClusterRoleBinding
  - **PATCH:** partially update the specified ClusterRoleBinding
  - **PUT:** replace the specified ClusterRoleBinding
- **/apis/rbac.authorization.k8s.io/v1/watch/clusterrolebindings/{name}**
  - **GET:** watch changes to an object of kind ClusterRoleBinding. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 21.3.2.1. /apis/rbac.authorization.k8s.io/v1/clusterrolebindings

Table 21.27. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**DELETE**

Description

delete collection of ClusterRoleBinding

Table 21.28. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 21.29. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.30. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ClusterRoleBinding

Table 21.31. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 21.32. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ClusterRoleBindingList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**



**Description**

create a ClusterRoleBinding

**Table 21.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 21.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ClusterRoleBinding</b> schema	

**Table 21.35. HTTP responses**

HTTP code	Response body
200 - OK	<b>ClusterRoleBinding</b> schema

HTTP code	Response body
201 - Created	<b>ClusterRoleBinding</b> schema
202 - Accepted	<b>ClusterRoleBinding</b> schema
401 - Unauthorized	Empty

### 21.3.2.2. /apis/rbac.authorization.k8s.io/v1/watch/clusterrolebindings

Table 21.36. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of ClusterRoleBinding. deprecated: use the 'watch' parameter with a list operation instead.

Table 21.37. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 21.3.2.3. /apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}

Table 21.38. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ClusterRoleBinding

Table 21.39. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a ClusterRoleBinding

Table 21.40. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 21.41. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.42. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ClusterRoleBinding

Table 21.43. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRoleBinding</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ClusterRoleBinding

**Table 21.44. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 21.45. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 21.46. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRoleBinding</b> schema
201 - Created	<b>ClusterRoleBinding</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ClusterRoleBinding

Table 21.47. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 21.48. Body parameters

Parameter	Type	Description
<b>body</b>	<b>ClusterRoleBinding</b> schema	

Table 21.49. HTTP responses

HTTP code	Response body
200 - OK	<b>ClusterRoleBinding</b> schema
201 - Created	<b>ClusterRoleBinding</b> schema
401 - Unauthorized	Empty

#### 21.3.2.4. /apis/rbac.authorization.k8s.io/v1/watch/clusterrolebindings/{name}

Table 21.50. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ClusterRoleBinding

Table 21.51. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ClusterRoleBinding. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 21.52. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 21.4. ROLE [RBAC.AUTHORIZATION.K8S.IO/V1]

### Description

Role is a namespaced, logical grouping of PolicyRules that can be referenced as a unit by a RoleBinding.

### Type

**object**

### 21.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata.
<b>rules</b>	<b>array</b>	Rules holds all the PolicyRules for this Role

Property	Type	Description
<b>rules[]</b>	<b>object</b>	PolicyRule holds information that describes a policy rule, but does not contain information about who the rule applies to or which namespace the rule applies to.

#### 21.4.1.1. .rules

##### Description

Rules holds all the PolicyRules for this Role

##### Type

**array**

#### 21.4.1.2. .rules[]

##### Description

PolicyRule holds information that describes a policy rule, but does not contain information about who the rule applies to or which namespace the rule applies to.

##### Type

**object**

##### Required

- **verbs**

Property	Type	Description
<b>apiGroups</b>	<b>array (string)</b>	APIGroups is the name of the APIGroup that contains the resources. If multiple API groups are specified, any action requested against one of the enumerated resources in any API group will be allowed. "" represents the core API group and "*" represents all API groups.

Property	Type	Description
<b>nonResourceURLs</b>	<b>array (string)</b>	NonResourceURLs is a set of partial urls that a user should have access to. *s are allowed, but only as the full, final step in the path. Since non-resource URLs are not namespaced, this field is only applicable for ClusterRoles referenced from a ClusterRoleBinding. Rules can either apply to API resources (such as "pods" or "secrets") or non-resource URL paths (such as "/api"), but not both.
<b>resourceNames</b>	<b>array (string)</b>	ResourceNames is an optional white list of names that the rule applies to. An empty set means that everything is allowed.
<b>resources</b>	<b>array (string)</b>	Resources is a list of resources this rule applies to. '*' represents all resources.
<b>verbs</b>	<b>array (string)</b>	Verbs is a list of Verbs that apply to ALL the ResourceKinds contained in this rule. '*' represents all verbs.

## 21.4.2. API endpoints

The following API endpoints are available:

- **/apis/rbac.authorization.k8s.io/v1/roles**
  - **GET**: list or watch objects of kind Role
- **/apis/rbac.authorization.k8s.io/v1/watch/roles**
  - **GET**: watch individual changes to a list of Role. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles**
  - **DELETE**: delete collection of Role
  - **GET**: list or watch objects of kind Role
  - **POST**: create a Role
- **/apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/roles**
  - **GET**: watch individual changes to a list of Role. deprecated: use the 'watch' parameter with

a list operation instead.

- **/apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}**
  - **DELETE:** delete a Role
  - **GET:** read the specified Role
  - **PATCH:** partially update the specified Role
  - **PUT:** replace the specified Role
- **/apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/roles/{name}**
  - **GET:** watch changes to an object of kind Role. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 21.4.2.1. /apis/rbac.authorization.k8s.io/v1/roles

Table 21.53. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind Role

**Table 21.54. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>RoleList</b> schema

HTTP code	Reponse body
401 - Unauthorized	Empty

### 21.4.2.2. /apis/rbac.authorization.k8s.io/v1/watch/roles

Table 21.55. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Role. deprecated: use the 'watch' parameter with a list operation instead.

**Table 21.56. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 21.4.2.3. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles

Table 21.57. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.58. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Role

Table 21.59. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 21.60. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.61. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind Role

Table 21.62. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 21.63. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">RoleList</a> schema
401 - Unauthorized	Empty

HTTP method  
**POST**

**Description**

create a Role

**Table 21.64. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 21.65. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Role</b> schema	

**Table 21.66. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>Role</b> schema

HTTP code	Response body
201 - Created	<a href="#">Role</a> schema
202 - Accepted	<a href="#">Role</a> schema
401 - Unauthorized	Empty

#### 21.4.2.4. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/roles

Table 21.67. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.68. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of Role. deprecated: use the 'watch' parameter with a list operation instead.

Table 21.69. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 21.4.2.5. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}

Table 21.70. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Role
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.71. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a Role

Table 21.72. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 21.73. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.74. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Role

Table 21.75. HTTP responses

HTTP code	Response body
200 - OK	<b>Role</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Role

**Table 21.76. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 21.77. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 21.78. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Role</b> schema
201 - Created	<b>Role</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Role

Table 21.79. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 21.80. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Role</b> schema	

Table 21.81. HTTP responses

HTTP code	Response body
200 - OK	<b>Role</b> schema
201 - Created	<b>Role</b> schema
401 - Unauthorized	Empty

#### 21.4.2.6. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/roles/{name}

Table 21.82. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Role
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects



Table 21.83. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind Role. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 21.84. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 21.5. ROLEBINDING [RBAC.AUTHORIZATION.K8S.IO/V1]

### Description

RoleBinding references a role, but does not contain it. It can reference a Role in the same namespace or a ClusterRole in the global namespace. It adds who information via Subjects and namespace information by which namespace it exists in. RoleBindings in a given namespace only have effect in that namespace.

### Type

**object**

### Required

- **roleRef**

### 21.5.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata.

Property	Type	Description
<b>roleRef</b>	<b>object</b>	RoleRef contains information that points to the role being used
<b>subjects</b>	<b>array</b>	Subjects holds references to the objects the role applies to.
<b>subjects[]</b>	<b>object</b>	Subject contains a reference to the object or user identities a role binding applies to. This can either hold a direct API object reference, or a value for non-objects such as user and group names.

### 21.5.1.1. .roleRef

#### Description

RoleRef contains information that points to the role being used

#### Type

**object**

#### Required

- **apiGroup**
- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup is the group for the resource being referenced
<b>kind</b>	<b>string</b>	Kind is the type of resource being referenced
<b>name</b>	<b>string</b>	Name is the name of resource being referenced

### 21.5.1.2. .subjects

#### Description

Subjects holds references to the objects the role applies to.

#### Type

**array**

### 21.5.1.3. .subjects[]

#### Description

Subject contains a reference to the object or user identities a role binding applies to. This can either hold a direct API object reference, or a value for non-objects such as user and group names.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
<b>apiGroup</b>	<b>string</b>	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
<b>kind</b>	<b>string</b>	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
<b>name</b>	<b>string</b>	Name of the object being referenced.
<b>namespace</b>	<b>string</b>	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

### 21.5.2. API endpoints

The following API endpoints are available:

- **/apis/rbac.authorization.k8s.io/v1/rolebindings**
  - **GET**: list or watch objects of kind RoleBinding
- **/apis/rbac.authorization.k8s.io/v1/watch/rolebindings**
  - **GET**: watch individual changes to a list of RoleBinding. deprecated: use the 'watch' parameter with a list operation instead.

- **/apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings**
  - **DELETE**: delete collection of RoleBinding
  - **GET**: list or watch objects of kind RoleBinding
  - **POST**: create a RoleBinding
- **/apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/rolebindings**
  - **GET**: watch individual changes to a list of RoleBinding. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}**
  - **DELETE**: delete a RoleBinding
  - **GET**: read the specified RoleBinding
  - **PATCH**: partially update the specified RoleBinding
  - **PUT**: replace the specified RoleBinding
- **/apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/rolebindings/{name}**
  - **GET**: watch changes to an object of kind RoleBinding. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 21.5.2.1. /apis/rbac.authorization.k8s.io/v1/rolebindings

Table 21.85. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind RoleBinding

**Table 21.86. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">RoleBindingList</a> schema

HTTP code	Reponse body
401 - Unauthorized	Empty

### 21.5.2.2. /apis/rbac.authorization.k8s.io/v1/watch/rolebindings

Table 21.87. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of RoleBinding. deprecated: use the 'watch' parameter with a list operation instead.

**Table 21.88. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 21.5.2.3. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings

Table 21.89. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.90. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of RoleBinding

Table 21.91. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>



Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 21.92. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.93. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind RoleBinding

Table 21.94. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>"k8s.io/initial-events-end": "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 21.95. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">RoleBindingList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a RoleBinding

**Table 21.96. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 21.97. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>RoleBinding</b> schema	

**Table 21.98. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>RoleBinding</b> schema

HTTP code	Reponse body
201 - Created	<a href="#">RoleBinding</a> schema
202 - Accepted	<a href="#">RoleBinding</a> schema
401 - Unauthorized	Empty

#### 21.5.2.4. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/rolebindings

Table 21.99. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.100. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of RoleBinding. deprecated: use the 'watch' parameter with a list operation instead.

Table 21.101. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 21.5.2.5. /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}

Table 21.102. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the RoleBinding
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.103. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a RoleBinding

Table 21.104. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 21.105. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 21.106. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified RoleBinding

Table 21.107. HTTP responses

HTTP code	Response body
200 - OK	<b>RoleBinding</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified RoleBinding

**Table 21.108. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 21.109. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 21.110. HTTP responses

HTTP code	Response body
200 - OK	<b>RoleBinding</b> schema
201 - Created	<b>RoleBinding</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified RoleBinding

Table 21.111. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 21.112. Body parameters

Parameter	Type	Description
<b>body</b>	<b>RoleBinding</b> schema	

Table 21.113. HTTP responses

HTTP code	Response body
200 - OK	<b>RoleBinding</b> schema
201 - Created	<b>RoleBinding</b> schema
401 - Unauthorized	Empty

### 21.5.2.6. /apis/rbac.authorization.k8s.io/v1/watch/namespaces/{namespace}/rolebindings/{name}

Table 21.114. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the RoleBinding
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 21.115. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind RoleBinding. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 21.116. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>WatchEvent</b> schema
401 - Unauthorized	Empty

## CHAPTER 22. NETWORK APIS

### 22.1. NETWORK APIS

#### 22.1.1. Route [route.openshift.io/v1]

##### Description

A route allows developers to expose services through an HTTP(S) aware load balancing and proxy layer via a public DNS entry. The route may further specify TLS options and a certificate, or specify a public CNAME that the router should also accept for HTTP and HTTPS traffic. An administrator typically configures their router to be visible outside the cluster firewall, and may also add additional security, caching, or traffic controls on the service content. Routers usually talk directly to the service endpoints. Once a route is created, the **host** field may not be changed. Generally, routers use the oldest route with a given host when resolving conflicts. Routers are subject to additional customization and may support additional controls via the annotations field. Because administrators may configure multiple routers, the route status field is used to return information to clients about the names and states of the route under each router. If a client chooses a duplicate name, for instance, the route status conditions are used to indicate the route cannot be chosen. To enable HTTP/2 ALPN on a route it requires a custom (non-wildcard) certificate. This prevents connection coalescing by clients, notably web browsers. We do not support HTTP/2 ALPN on routes that use the default certificate because of the risk of connection re-use/coalescing. Routes that do not have their own custom certificate will not be HTTP/2 ALPN-enabled on either the frontend or the backend. Compatibility level 1: Stable within a major release for a minimum of 12 months or 3 minor releases (whichever is longer).

##### Type

**object**

### 22.2. ROUTE [ROUTE.OPENSIFT.IO/V1]

##### Description

A route allows developers to expose services through an HTTP(S) aware load balancing and proxy layer via a public DNS entry. The route may further specify TLS options and a certificate, or specify a public CNAME that the router should also accept for HTTP and HTTPS traffic. An administrator typically configures their router to be visible outside the cluster firewall, and may also add additional security, caching, or traffic controls on the service content. Routers usually talk directly to the service endpoints. Once a route is created, the **host** field may not be changed. Generally, routers use the oldest route with a given host when resolving conflicts. Routers are subject to additional customization and may support additional controls via the annotations field. Because administrators may configure multiple routers, the route status field is used to return information to clients about the names and states of the route under each router. If a client chooses a duplicate name, for instance, the route status conditions are used to indicate the route cannot be chosen. To enable HTTP/2 ALPN on a route it requires a custom (non-wildcard) certificate. This prevents connection coalescing by clients, notably web browsers. We do not support HTTP/2 ALPN on routes that use the default certificate because of the risk of connection re-use/coalescing. Routes that do not have their own custom certificate will not be HTTP/2 ALPN-enabled on either the frontend or the backend. Compatibility level 1: Stable within a major release for a minimum of 12 months or 3 minor releases (whichever is longer).

##### Type

**object**

##### Required

- **spec**

## 22.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	spec is the desired state of the route
<b>status</b>	<b>object</b>	status is the current state of the route

## 22.2.1.1. .spec

## Description

spec is the desired state of the route

## Type

**object**

## Required

- **to**

Property	Type	Description
<b>alternateBackends</b>	<b>array</b>	alternateBackends allows up to 3 additional backends to be assigned to the route. Only the Service kind is allowed, and it will be defaulted to Service. Use the weight field in RouteTargetReference object to specify relative preference.
<b>alternateBackends[]</b>	<b>object</b>	RouteTargetReference specifies the target that resolve into endpoints. Only the 'Service' kind is allowed. Use 'weight' field to emphasize one over others.
<b>host</b>	<b>string</b>	host is an alias/DNS that points to the service. Optional. If not specified a route name will typically be automatically chosen. Must follow DNS952 subdomain conventions.
<b>path</b>	<b>string</b>	path that the router watches for, to route traffic for to the service. Optional
<b>port</b>	<b>object</b>	If specified, the port to be used by the router. Most routers will use all endpoints exposed by the service by default - set this value to instruct routers which port to use.

Property	Type	Description
<b>subdomain</b>	<b>string</b>	subdomain is a DNS subdomain that is requested within the ingress controller's domain (as a subdomain). If host is set this field is ignored. An ingress controller may choose to ignore this suggested name, in which case the controller will report the assigned name in the status.ingress array or refuse to admit the route. If this value is set and the server does not support this field host will be populated automatically. Otherwise host is left empty. The field may have multiple parts separated by a dot, but not all ingress controllers may honor the request. This field may not be changed after creation except by a user with the update routes/custom-host permission. Example: subdomain <b>frontend</b> automatically receives the router subdomain <b>apps.mycluster.com</b> to have a full hostname <b>frontend.apps.mycluster.com</b> .
<b>tls</b>	<b>object</b>	The tls field provides the ability to configure certificates and termination for the route.
<b>to</b>	<b>object</b>	to is an object the route should use as the primary backend. Only the Service kind is allowed, and it will be defaulted to Service. If the weight field (0-256 default 100) is set to zero, no traffic will be sent to this backend.
<b>wildcardPolicy</b>	<b>string</b>	Wildcard policy if any for the route. Currently only 'Subdomain' or 'None' is allowed.

### 22.2.1.2. .spec.alternateBackends

#### Description

alternateBackends allows up to 3 additional backends to be assigned to the route. Only the Service kind is allowed, and it will be defaulted to Service. Use the weight field in RouteTargetReference object to specify relative preference.

#### Type

**array**

### 22.2.1.3. .spec.alternateBackends[]

#### Description

RouteTargetReference specifies the target that resolve into endpoints. Only the 'Service' kind is allowed. Use 'weight' field to emphasize one over others.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
<b>kind</b>	<b>string</b>	The kind of target that the route is referring to. Currently, only 'Service' is allowed
<b>name</b>	<b>string</b>	name of the service/target that is being referred to. e.g. name of the service
<b>weight</b>	<b>integer</b>	weight as an integer between 0 and 256, default 100, that specifies the target's relative weight against other target reference objects. 0 suppresses requests to this backend.

### 22.2.1.4. .spec.port

#### Description

If specified, the port to be used by the router. Most routers will use all endpoints exposed by the service by default - set this value to instruct routers which port to use.

#### Type

**object**

#### Required

- **targetPort**



Property	Type	Description
<b>targetPort</b>	<b>integer-or-string</b>	

### 22.2.1.5. .spec.tls

#### Description

The `tls` field provides the ability to configure certificates and termination for the route.

#### Type

**object**

#### Required

- **termination**

Property	Type	Description
<b>caCertificate</b>	<b>string</b>	<code>caCertificate</code> provides the cert authority certificate contents
<b>certificate</b>	<b>string</b>	<code>certificate</code> provides certificate contents. This should be a single serving certificate, not a certificate chain. Do not include a CA certificate.
<b>destinationCACertificate</b>	<b>string</b>	<code>destinationCACertificate</code> provides the contents of the ca certificate of the final destination. When using reencrypt termination this file should be provided in order to have routers use it for health checks on the secure connection. If this field is not specified, the router may provide its own destination CA and perform hostname validation using the short service name ( <code>service.namespace.svc</code> ), which allows infrastructure generated certificates to automatically verify.

Property	Type	Description
<b>insecureEdgeTerminationPolicy</b>	<b>string</b>	insecureEdgeTerminationPolicy indicates the desired behavior for insecure connections to a route. While each router may make its own decisions on which ports to expose, this is normally port 80. * Allow - traffic is sent to the server on the insecure port (edge/reencrypt terminations only) (default). * None - no traffic is allowed on the insecure port. * Redirect - clients are redirected to the secure port.
<b>key</b>	<b>string</b>	key provides key file contents
<b>termination</b>	<b>string</b>	termination indicates termination type. * edge - TLS termination is done by the router and http is used to communicate with the backend (default) * passthrough - Traffic is sent straight to the destination without the router providing TLS termination * reencrypt - TLS termination is done by the router and https is used to communicate with the backend

### 22.2.1.6. .spec.to

#### Description

to is an object the route should use as the primary backend. Only the Service kind is allowed, and it will be defaulted to Service. If the weight field (0-256 default 100) is set to zero, no traffic will be sent to this backend.

#### Type

**object**

#### Required

- **kind**
- **name**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>kind</b>	<b>string</b>	The kind of target that the route is referring to. Currently, only 'Service' is allowed
<b>name</b>	<b>string</b>	name of the service/target that is being referred to. e.g. name of the service
<b>weight</b>	<b>integer</b>	weight as an integer between 0 and 256, default 100, that specifies the target's relative weight against other target reference objects. 0 suppresses requests to this backend.

### 22.2.1.7. .status

#### Description

status is the current state of the route

#### Type

**object**

Property	Type	Description
<b>ingress</b>	<b>array</b>	ingress describes the places where the route may be exposed. The list of ingress points may contain duplicate Host or RouterName values. Routes are considered live once they are <b>Ready</b>
<b>ingress[]</b>	<b>object</b>	RouteIngress holds information about the places where a route is exposed.

### 22.2.1.8. .status.ingress

#### Description

ingress describes the places where the route may be exposed. The list of ingress points may contain duplicate Host or RouterName values. Routes are considered live once they are **Ready**

#### Type

**array**

### 22.2.1.9. .status.ingress[]

#### Description

RouteIngress holds information about the places where a route is exposed.

## Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array</b>	Conditions is the state of the route, may be empty.
<b>conditions[]</b>	<b>object</b>	RouteIngressCondition contains details for the current condition of this route on a particular router.
<b>host</b>	<b>string</b>	Host is the host string under which the route is exposed; this value is required
<b>routerCanonicalHostname</b>	<b>string</b>	CanonicalHostname is the external host name for the router that can be used as a CNAME for the host requested for this route. This value is optional and may not be set in all cases.
<b>routerName</b>	<b>string</b>	Name is a name chosen by the router to identify itself; this value is required
<b>wildcardPolicy</b>	<b>string</b>	Wildcard policy is the wildcard policy that was allowed where this route is exposed.

### 22.2.1.10. .status.ingress[].conditions

#### Description

Conditions is the state of the route, may be empty.

#### Type

**array**

### 22.2.1.11. .status.ingress[].conditions[]

#### Description

RouteIngressCondition contains details for the current condition of this route on a particular router.

#### Type

**object**

#### Required

- **status**

- **type**

Property	Type	Description
<b>lastTransitionTime</b>	<b>string</b>	RFC 3339 date and time when this condition last transitioned
<b>message</b>	<b>string</b>	Human readable message indicating details about last transition.
<b>reason</b>	<b>string</b>	(brief) reason for the condition's last transition, and is usually a machine and human readable constant
<b>status</b>	<b>string</b>	Status is the status of the condition. Can be True, False, Unknown.
<b>type</b>	<b>string</b>	Type is the type of the condition. Currently only Admitted.

### 22.2.2. API endpoints

The following API endpoints are available:

- **/apis/route.openshift.io/v1/routes**
  - **GET**: list objects of kind Route
- **/apis/route.openshift.io/v1/namespaces/{namespace}/routes**
  - **DELETE**: delete collection of Route
  - **GET**: list objects of kind Route
  - **POST**: create a Route
- **/apis/route.openshift.io/v1/namespaces/{namespace}/routes/{name}**
  - **DELETE**: delete a Route
  - **GET**: read the specified Route
  - **PATCH**: partially update the specified Route
  - **PUT**: replace the specified Route
- **/apis/route.openshift.io/v1/namespaces/{namespace}/routes/{name}/status**
  - **GET**: read status of the specified Route

- **PATCH**: partially update status of the specified Route
- **PUT**: replace status of the specified Route

### 22.2.2.1. /apis/route.openshift.io/v1/routes

Table 22.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list objects of kind Route

Table 22.2. HTTP responses



HTTP code	Reponse body
200 - OK	<a href="#">RouteList</a> schema
401 - Unauthorized	Empty

### 22.2.2.2. /apis/route.openshift.io/v1/namespaces/{namespace}/routes

Table 22.3. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 22.4. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of Route

Table 22.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 22.6. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

HTTP method

**GET**

**Description**

list objects of kind Route

**Table 22.7. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 22.8. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">RouteList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a Route

**Table 22.9. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 22.10. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>Route</b> schema	

**Table 22.11. HTTP responses**

HTTP code	Response body
200 - OK	<b>Route</b> schema



HTTP code	Response body
201 - Created	<a href="#">Route</a> schema
202 - Accepted	<a href="#">Route</a> schema
401 - Unauthorized	Empty

### 22.2.2.3. /apis/route.openshift.io/v1/namespaces/{namespace}/routes/{name}

Table 22.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Route
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 22.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a Route

Table 22.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 22.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 22.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified Route

Table 22.17. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 22.18. HTTP responses

HTTP code	Response body
200 - OK	<b>Route</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified Route

Table 22.19. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 22.20. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 22.21. HTTP responses

HTTP code	Response body
200 - OK	<b>Route</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified Route

Table 22.22. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 22.23. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Route</b> schema	

Table 22.24. HTTP responses

HTTP code	Response body
200 - OK	<b>Route</b> schema
201 - Created	<b>Route</b> schema
401 - Unauthorized	Empty

## 22.2.2.4. /apis/route.openshift.io/v1/namespaces/{namespace}/routes/{name}/status

Table 22.25. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the Route
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 22.26. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

read status of the specified Route

Table 22.27. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 22.28. HTTP responses

HTTP code	Response body
200 - OK	<b>Route</b> schema
401 - Unauthorized	Empty

## HTTP method

**PATCH**

## Description

partially update status of the specified Route

Table 22.29. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 22.30. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 22.31. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">Route</a> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified Route

**Table 22.32. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 22.33. Body parameters**



Parameter	Type	Description
<b>body</b>	<b>Route</b> schema	

Table 22.34. HTTP responses

HTTP code	Response body
200 - OK	<b>Route</b> schema
201 - Created	<b>Route</b> schema
401 - Unauthorized	Empty

## CHAPTER 23. SCHEDULING APIS

### 23.1. SCHEDULING APIS

#### 23.1.1. PriorityClass [scheduling.k8s.io/v1]

##### Description

PriorityClass defines mapping from a priority class name to the priority integer value. The value can be any valid integer.

##### Type

**object**

### 23.2. PRIORITYCLASS [SCHEDULING.K8S.IO/V1]

##### Description

PriorityClass defines mapping from a priority class name to the priority integer value. The value can be any valid integer.

##### Type

**object**

##### Required

- **value**

#### 23.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>description</b>	<b>string</b>	description is an arbitrary string that usually provides guidelines on when this priority class should be used.

Property	Type	Description
<b>globalDefault</b>	<b>boolean</b>	globalDefault specifies whether this PriorityClass should be considered as the default priority for pods that do not have any priority class. Only one PriorityClass can be marked as <b>globalDefault</b> . However, if more than one PriorityClasses exists with their <b>globalDefault</b> field set to true, the smallest value of such global default PriorityClasses will be used as the default priority.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>preemptionPolicy</b>	<b>string</b>	preemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset.  Possible enum values: - <b>"Never"</b> means that pod never preempts other pods with lower priority. - <b>"PreemptLowerPriority"</b> means that pod can preempt other pods with lower priority.
<b>value</b>	<b>integer</b>	value represents the integer value of this priority class. This is the actual priority that pods receive when they have the name of this class in their pod spec.

## 23.2.2. API endpoints

The following API endpoints are available:

- **/apis/scheduling.k8s.io/v1/priorityclasses**
  - **DELETE**: delete collection of PriorityClass
  - **GET**: list or watch objects of kind PriorityClass
  - **POST**: create a PriorityClass
- **/apis/scheduling.k8s.io/v1/watch/priorityclasses**
  - **GET**: watch individual changes to a list of PriorityClass. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/scheduling.k8s.io/v1/priorityclasses/{name}**
  - **DELETE**: delete a PriorityClass
  - **GET**: read the specified PriorityClass
  - **PATCH**: partially update the specified PriorityClass
  - **PUT**: replace the specified PriorityClass
- **/apis/scheduling.k8s.io/v1/watch/priorityclasses/{name}**
  - **GET**: watch changes to an object of kind PriorityClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 23.2.2.1. /apis/scheduling.k8s.io/v1/priorityclasses

Table 23.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**DELETE**

Description

delete collection of PriorityClass

Table 23.2. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 23.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 23.4. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind PriorityClass

Table 23.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 23.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">PriorityClassList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a PriorityClass

**Table 23.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 23.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>PriorityClass</b> schema	

**Table 23.9. HTTP responses**

HTTP code	Response body
200 - OK	<b>PriorityClass</b> schema

HTTP code	Response body
201 - Created	<b>PriorityClass</b> schema
202 - Accepted	<b>PriorityClass</b> schema
401 - Unauthorized	Empty

### 23.2.2.2. /apis/scheduling.k8s.io/v1/watch/priorityclasses

Table 23.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
GET

## Description

watch individual changes to a list of PriorityClass. deprecated: use the 'watch' parameter with a list operation instead.

**Table 23.11. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 23.2.2.3. /apis/scheduling.k8s.io/v1/priorityclasses/{name}

**Table 23.12. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PriorityClass

**Table 23.13. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

### DELETE

## Description

delete a PriorityClass

**Table 23.14. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 23.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 23.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified PriorityClass

Table 23.17. HTTP responses

HTTP code	Response body
200 - OK	<b>PriorityClass</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified PriorityClass

**Table 23.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 23.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 23.20. HTTP responses

HTTP code	Reponse body
200 - OK	<b>PriorityClass</b> schema
201 - Created	<b>PriorityClass</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified PriorityClass

Table 23.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 23.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>PriorityClass</b> schema	

Table 23.23. HTTP responses

HTTP code	Response body
200 - OK	<b>PriorityClass</b> schema
201 - Created	<b>PriorityClass</b> schema
401 - Unauthorized	Empty

#### 23.2.2.4. /apis/scheduling.k8s.io/v1/watch/priorityclasses/{name}

Table 23.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the PriorityClass

Table 23.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind PriorityClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 23.26. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty



## CHAPTER 24. SECURITY APIS

### 24.1. SECURITY APIS

#### 24.1.1. SecurityContextConstraints [security.openshift.io/v1]

##### Description

SecurityContextConstraints governs the ability to make requests that affect the SecurityContext that will be applied to a container. For historical reasons SCC was exposed under the core Kubernetes API group. That exposure is deprecated and will be removed in a future release - users should instead use the security.openshift.io group to manage SecurityContextConstraints. Compatibility level 1: Stable within a major release for a minimum of 12 months or 3 minor releases (whichever is longer).

##### Type

**object**

### 24.2. SECURITYCONTEXTCONSTRAINTS [SECURITY.OPENSIFT.IO/V1]

##### Description

SecurityContextConstraints governs the ability to make requests that affect the SecurityContext that will be applied to a container. For historical reasons SCC was exposed under the core Kubernetes API group. That exposure is deprecated and will be removed in a future release - users should instead use the security.openshift.io group to manage SecurityContextConstraints. Compatibility level 1: Stable within a major release for a minimum of 12 months or 3 minor releases (whichever is longer).

##### Type

**object**

##### Required

- **allowHostDirVolumePlugin**
- **allowHostIPC**
- **allowHostNetwork**
- **allowHostPID**
- **allowHostPorts**
- **allowPrivilegedContainer**
- **readOnlyRootFilesystem**

#### 24.2.1. Specification

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>allowHostDirVolumePlugin</b>	<b>boolean</b>	AllowHostDirVolumePlugin determines if the policy allow containers to use the HostDir volume plugin
<b>allowHostIPC</b>	<b>boolean</b>	AllowHostIPC determines if the policy allows host ipc in the containers.
<b>allowHostNetwork</b>	<b>boolean</b>	AllowHostNetwork determines if the policy allows the use of HostNetwork in the pod spec.
<b>allowHostPID</b>	<b>boolean</b>	AllowHostPID determines if the policy allows host pid in the containers.
<b>allowHostPorts</b>	<b>boolean</b>	AllowHostPorts determines if the policy allows host ports in the containers.
<b>allowPrivilegeEscalation</b>	``	AllowPrivilegeEscalation determines if a pod can request to allow privilege escalation. If unspecified, defaults to true.
<b>allowPrivilegedContainer</b>	<b>boolean</b>	AllowPrivilegedContainer determines if a container can request to be run as privileged.
<b>allowedCapabilities</b>	``	AllowedCapabilities is a list of capabilities that can be requested to add to the container. Capabilities in this field maybe added at the pod author's discretion. You must not list a capability in both AllowedCapabilities and RequiredDropCapabilities. To allow all capabilities you may use '*'.

Property	Type	Description
<b>allowedFlexVolumes</b>	``	AllowedFlexVolumes is a whitelist of allowed Flexvolumes. Empty or nil indicates that all Flexvolumes may be used. This parameter is effective only when the usage of the Flexvolumes is allowed in the "Volumes" field.
<b>allowedUnsafeSysctls</b>	``	AllowedUnsafeSysctls is a list of explicitly allowed unsafe sysctls, defaults to none. Each entry is either a plain sysctl name or ends in "" in which case it is considered as a prefix of allowed sysctls. Single * means all unsafe sysctls are allowed. Kubelet has to whitelist all allowed unsafe sysctls explicitly to avoid rejection. Examples: e.g. "foo/" allows "foo/bar", "foo/baz", etc. e.g. "foo.*" allows "foo.bar", "foo.baz", etc.
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>defaultAddCapabilities</b>	``	DefaultAddCapabilities is the default set of capabilities that will be added to the container unless the pod spec specifically drops the capability. You may not list a capability in both DefaultAddCapabilities and RequiredDropCapabilities.
<b>defaultAllowPrivilegeEscalation</b>	``	DefaultAllowPrivilegeEscalation controls the default setting for whether a process can gain more privileges than its parent process.

Property	Type	Description
<b>forbiddenSysctls</b>	``	ForbiddenSysctls is a list of explicitly forbidden sysctls, defaults to none. Each entry is either a plain sysctl name or ends in "" in which case it is considered as a prefix of forbidden sysctls. Single * means all sysctls are forbidden. Examples: e.g. "foo/" forbids "foo/bar", "foo/baz", etc. e.g. "foo.*" forbids "foo.bar", "foo.baz", etc.
<b>fsGroup</b>	``	FSGroup is the strategy that will dictate what fs group is used by the SecurityContext.
<b>groups</b>	``	The groups that have permission to use this security context constraints
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>priority</b>	``	Priority influences the sort order of SCCs when evaluating which SCCs to try first for a given pod request based on access in the Users and Groups fields. The higher the int, the higher priority. An unset value is considered a 0 priority. If scores for multiple SCCs are equal they will be sorted from most restrictive to least restrictive. If both priorities and restrictions are equal the SCCs will be sorted by name.
<b>readOnlyRootFilesystem</b>	<b>boolean</b>	ReadOnlyRootFilesystem when set to true will force containers to run with a read only root file system. If the container specifically requests to run with a non-read only root file system the SCC should deny the pod. If set to false the container may run with a read only root file system if it wishes but it will not be forced to.
<b>requiredDropCapabilities</b>	``	RequiredDropCapabilities are the capabilities that will be dropped from the container. These are required to be dropped and cannot be added.
<b>runAsUser</b>	``	RunAsUser is the strategy that will dictate what RunAsUser is used in the SecurityContext.
<b>seLinuxContext</b>	``	SELinuxContext is the strategy that will dictate what labels will be set in the SecurityContext.
<b>seccompProfiles</b>	``	SeccompProfiles lists the allowed profiles that may be set for the pod or container's seccomp annotations. An unset (nil) or empty value means that no profiles may be specified by the pod or container. The wildcard '*' may be used to allow all profiles. When used to generate a value for a pod the first non-wildcard profile will be used as the default.

Property	Type	Description
<b>supplementalGroups</b>	``	SupplementalGroups is the strategy that will dictate what supplemental groups are used by the SecurityContext.
<b>users</b>	``	The users who have permissions to use this security context constraints
<b>volumes</b>	``	Volumes is a white list of allowed volume plugins. FSType corresponds directly with the field names of a VolumeSource (azureFile, configMap, emptyDir). To allow all volumes you may use "*". To allow no volumes, set to ["none"].

## 24.2.2. API endpoints

The following API endpoints are available:

- **/apis/security.openshift.io/v1/securitycontextconstraints**
  - **DELETE**: delete collection of SecurityContextConstraints
  - **GET**: list objects of kind SecurityContextConstraints
  - **POST**: create SecurityContextConstraints
- **/apis/security.openshift.io/v1/securitycontextconstraints/{name}**
  - **DELETE**: delete SecurityContextConstraints
  - **GET**: read the specified SecurityContextConstraints
  - **PATCH**: partially update the specified SecurityContextConstraints
  - **PUT**: replace the specified SecurityContextConstraints

### 24.2.2.1. /apis/security.openshift.io/v1/securitycontextconstraints

Table 24.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**DELETE**

**Description**

delete collection of SecurityContextConstraints

**Table 24.2. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 24.3. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

## HTTP method

**GET****Description**

list objects of kind SecurityContextConstraints

**Table 24.4. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 24.5. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">SecurityContextConstraintsList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

Description

create SecurityContextConstraints

Table 24.6. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 24.7. Body parameters

Parameter	Type	Description
<b>body</b>	<b>SecurityContextConstraints</b> schema	

Table 24.8. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">SecurityContextConstraints</a> schema
201 - Created	<a href="#">SecurityContextConstraints</a> schema
202 - Accepted	<a href="#">SecurityContextConstraints</a> schema
401 - Unauthorized	Empty

#### 24.2.2.2. /apis/security.openshift.io/v1/securitycontextconstraints/{name}

Table 24.9. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the SecurityContextConstraints

Table 24.10. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### DELETE

#### Description

delete SecurityContextConstraints

Table 24.11. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 24.12. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 24.13. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified SecurityContextConstraints

Table 24.14. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 24.15. HTTP responses

HTTP code	Response body
200 - OK	<b>SecurityContextConstraints</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified SecurityContextConstraints

Table 24.16. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 24.17. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 24.18. HTTP responses

HTTP code	Response body
200 - OK	<b>SecurityContextConstraints</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified SecurityContextConstraints

Table 24.19. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 24.20. Body parameters

Parameter	Type	Description
<b>body</b>	<b>SecurityContextConstraints</b> schema	

Table 24.21. HTTP responses

HTTP code	Response body
200 - OK	<b>SecurityContextConstraints</b> schema
201 - Created	<b>SecurityContextConstraints</b> schema
401 - Unauthorized	Empty

## CHAPTER 25. SECURITY-INTERNAL APIS

### 25.1. SECURITY INTERNAL APIS

#### 25.1.1. RangeAllocation [security.internal.openshift.io/v1]

##### Description

RangeAllocation is used so we can easily expose a RangeAllocation typed for security group This is an internal API, not intended for external consumption. Compatibility level 1: Stable within a major release for a minimum of 12 months or 3 minor releases (whichever is longer).

##### Type

**object**

### 25.2. RANGEALLOCATION [SECURITY.INTERNAL.OPENSIFT.IO/V1]

##### Description

RangeAllocation is used so we can easily expose a RangeAllocation typed for security group This is an internal API, not intended for external consumption. Compatibility level 1: Stable within a major release for a minimum of 12 months or 3 minor releases (whichever is longer).

##### Type

**object**

#### 25.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>data</b>	<b>string</b>	data is a byte array representing the serialized state of a range allocation. It is a bitmap with each bit set to one to represent a range is taken.

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>range</b>	<b>string</b>	range is a string representing a unique label for a range of uids, "10000000000-20000000000/10000".

### 25.2.2. API endpoints

The following API endpoints are available:

- **/apis/security.internal.openshift.io/v1/rangeallocations**
  - **DELETE**: delete collection of RangeAllocation
  - **GET**: list objects of kind RangeAllocation
  - **POST**: create a RangeAllocation
- **/apis/security.internal.openshift.io/v1/rangeallocations/{name}**
  - **DELETE**: delete a RangeAllocation
  - **GET**: read the specified RangeAllocation
  - **PATCH**: partially update the specified RangeAllocation
  - **PUT**: replace the specified RangeAllocation

#### 25.2.2.1. /apis/security.internal.openshift.io/v1/rangeallocations

Table 25.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**DELETE**

## Description

delete collection of RangeAllocation

Table 25.2. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 25.3. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

## HTTP method

**GET****Description**

list objects of kind RangeAllocation

**Table 25.4. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 25.5. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">RangeAllocationList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

Description

create a RangeAllocation

Table 25.6. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 25.7. Body parameters

Parameter	Type	Description
<b>body</b>	<b>RangeAllocation</b> schema	

Table 25.8. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">RangeAllocation</a> schema
201 - Created	<a href="#">RangeAllocation</a> schema
202 - Accepted	<a href="#">RangeAllocation</a> schema
401 - Unauthorized	Empty

### 25.2.2.2. /apis/security.internal.openshift.io/v1/rangeallocations/{name}

Table 25.9. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the RangeAllocation

Table 25.10. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a RangeAllocation

Table 25.11. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 25.12. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 25.13. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified RangeAllocation

Table 25.14. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 25.15. HTTP responses

HTTP code	Response body
200 - OK	<b>RangeAllocation</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified RangeAllocation

Table 25.16. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 25.17. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 25.18. HTTP responses

HTTP code	Response body
200 - OK	<b>RangeAllocation</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified RangeAllocation

Table 25.19. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 25.20. Body parameters

Parameter	Type	Description
<b>body</b>	<b>RangeAllocation</b> schema	

Table 25.21. HTTP responses

HTTP code	Response body
200 - OK	<b>RangeAllocation</b> schema
201 - Created	<b>RangeAllocation</b> schema
401 - Unauthorized	Empty



## CHAPTER 26. SNAPSHOT APIS

### 26.1. SNAPSHOT APIS

#### 26.1.1. VolumeSnapshot [snapshot.storage.k8s.io/v1]

##### Description

VolumeSnapshot is a user's request for either creating a point-in-time snapshot of a persistent volume, or binding to a pre-existing snapshot.

##### Type

**object**

#### 26.1.2. VolumeSnapshotClass [snapshot.storage.k8s.io/v1]

##### Description

VolumeSnapshotClass specifies parameters that a underlying storage system uses when creating a volume snapshot. A specific VolumeSnapshotClass is used by specifying its name in a VolumeSnapshot object. VolumeSnapshotClasses are non-namespaced

##### Type

**object**

#### 26.1.3. VolumeSnapshotContent [snapshot.storage.k8s.io/v1]

##### Description

VolumeSnapshotContent represents the actual "on-disk" snapshot object in the underlying storage system

##### Type

**object**

### 26.2. VOLUMESNAPSHOT [SNAPSHOT.STORAGE.K8S.IO/V1]

##### Description

VolumeSnapshot is a user's request for either creating a point-in-time snapshot of a persistent volume, or binding to a pre-existing snapshot.

##### Type

**object**

##### Required

- **spec**

#### 26.2.1. Specification

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	spec defines the desired characteristics of a snapshot requested by a user. More info: <a href="https://kubernetes.io/docs/concepts/storage/volume-snapshots#volumesnapshots">https://kubernetes.io/docs/concepts/storage/volume-snapshots#volumesnapshots</a> Required.
<b>status</b>	<b>object</b>	status represents the current information of a snapshot. Consumers must verify binding between VolumeSnapshot and VolumeSnapshotContent objects is successful (by validating that both VolumeSnapshot and VolumeSnapshotContent point at each other) before using this object.

## 26.2.1.1. .spec

**Description**

spec defines the desired characteristics of a snapshot requested by a user. More info: <https://kubernetes.io/docs/concepts/storage/volume-snapshots#volumesnapshots> Required.

**Type**

**object**

**Required**

- **source**

Property	Type	Description
<b>source</b>	<b>object</b>	source specifies where a snapshot will be created from. This field is immutable after creation. Required.
<b>volumeSnapshotClassName</b>	<b>string</b>	VolumeSnapshotClassName is the name of the VolumeSnapshotClass requested by the VolumeSnapshot. VolumeSnapshotClassName may be left nil to indicate that the default SnapshotClass should be used. A given cluster may have multiple default Volume SnapshotClasses: one default per CSI Driver. If a VolumeSnapshot does not specify a SnapshotClass, VolumeSnapshotSource will be checked to figure out what the associated CSI Driver is, and the default VolumeSnapshotClass associated with that CSI Driver will be used. If more than one VolumeSnapshotClass exist for a given CSI Driver and more than one have been marked as default, CreateSnapshot will fail and generate an event. Empty string is not allowed for this field.

**26.2.1.2. .spec.source****Description**

source specifies where a snapshot will be created from. This field is immutable after creation. Required.

**Type**

**object**

Property	Type	Description
<b>persistentVolumeClaimName</b>	<b>string</b>	persistentVolumeClaimName specifies the name of the PersistentVolumeClaim object representing the volume from which a snapshot should be created. This PVC is assumed to be in the same namespace as the VolumeSnapshot object. This field should be set if the snapshot does not exist, and needs to be created. This field is immutable.
<b>volumeSnapshotContentName</b>	<b>string</b>	volumeSnapshotContentName specifies the name of a pre-existing VolumeSnapshotContent object representing an existing volume snapshot. This field should be set if the snapshot already exists and only needs a representation in Kubernetes. This field is immutable.

### 26.2.1.3. .status

#### Description

status represents the current information of a snapshot. Consumers must verify binding between VolumeSnapshot and VolumeSnapshotContent objects is successful (by validating that both VolumeSnapshot and VolumeSnapshotContent point at each other) before using this object.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>boundVolumeSnapshotContentName</b>	<b>string</b>	boundVolumeSnapshotContentName is the name of the VolumeSnapshotContent object to which this VolumeSnapshot object intends to bind to. If not specified, it indicates that the VolumeSnapshot object has not been successfully bound to a VolumeSnapshotContent object yet. NOTE: To avoid possible security issues, consumers must verify binding between VolumeSnapshot and VolumeSnapshotContent objects is successful (by validating that both VolumeSnapshot and VolumeSnapshotContent point at each other) before using this object.
<b>creationTime</b>	<b>string</b>	creationTime is the timestamp when the point-in-time snapshot is taken by the underlying storage system. In dynamic snapshot creation case, this field will be filled in by the snapshot controller with the "creation_time" value returned from CSI "CreateSnapshot" gRPC call. For a pre-existing snapshot, this field will be filled with the "creation_time" value returned from the CSI "ListSnapshots" gRPC call if the driver supports it. If not specified, it may indicate that the creation time of the snapshot is unknown.

Property	Type	Description
<b>error</b>	<b>object</b>	error is the last observed error during snapshot creation, if any. This field could be helpful to upper level controllers(i.e., application controller) to decide whether they should continue on waiting for the snapshot to be created based on the type of error reported. The snapshot controller will keep retrying when an error occurs during the snapshot creation. Upon success, this error field will be cleared.
<b>readyToUse</b>	<b>boolean</b>	readyToUse indicates if the snapshot is ready to be used to restore a volume. In dynamic snapshot creation case, this field will be filled in by the snapshot controller with the "ready_to_use" value returned from CSI "CreateSnapshot" gRPC call. For a pre-existing snapshot, this field will be filled with the "ready_to_use" value returned from the CSI "ListSnapshots" gRPC call if the driver supports it, otherwise, this field will be set to "True". If not specified, it means the readiness of a snapshot is unknown.

Property	Type	Description
<b>restoreSize</b>	<b>integer-or-string</b>	restoreSize represents the minimum size of volume required to create a volume from this snapshot. In dynamic snapshot creation case, this field will be filled in by the snapshot controller with the "size_bytes" value returned from CSI "CreateSnapshot" gRPC call. For a pre-existing snapshot, this field will be filled with the "size_bytes" value returned from the CSI "ListSnapshots" gRPC call if the driver supports it. When restoring a volume from this snapshot, the size of the volume MUST NOT be smaller than the restoreSize if it is specified, otherwise the restoration will fail. If not specified, it indicates that the size is unknown.

#### 26.2.1.4. .status.error

##### Description

error is the last observed error during snapshot creation, if any. This field could be helpful to upper level controllers(i.e., application controller) to decide whether they should continue on waiting for the snapshot to be created based on the type of error reported. The snapshot controller will keep retrying when an error occurs during the snapshot creation. Upon success, this error field will be cleared.

##### Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	message is a string detailing the encountered error during snapshot creation if specified. NOTE: message may be logged, and it should not contain sensitive information.
<b>time</b>	<b>string</b>	time is the timestamp when the error was encountered.

#### 26.2.2. API endpoints

The following API endpoints are available:

- **/apis/snapshot.storage.k8s.io/v1/volumesnapshots**
  - **GET**: list objects of kind VolumeSnapshot
- **/apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots**
  - **DELETE**: delete collection of VolumeSnapshot
  - **GET**: list objects of kind VolumeSnapshot
  - **POST**: create a VolumeSnapshot
- **/apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots/{name}**
  - **DELETE**: delete a VolumeSnapshot
  - **GET**: read the specified VolumeSnapshot
  - **PATCH**: partially update the specified VolumeSnapshot
  - **PUT**: replace the specified VolumeSnapshot
- **/apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots/{name}/status**
  - **GET**: read status of the specified VolumeSnapshot
  - **PATCH**: partially update status of the specified VolumeSnapshot
  - **PUT**: replace status of the specified VolumeSnapshot

### 26.2.2.1. /apis/snapshot.storage.k8s.io/v1/volumesnapshots

Table 26.1. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list objects of kind VolumeSnapshot

**Table 26.2. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>VolumeSnapshotList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 26.2.2.2. /apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots

Table 26.3. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 26.4. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of VolumeSnapshot

Table 26.5. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 26.6. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

HTTP method

**GET**

**Description**

list objects of kind VolumeSnapshot

**Table 26.7. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.



Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 26.8. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a VolumeSnapshot

**Table 26.9. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 26.10. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshot</b> schema	

**Table 26.11. HTTP responses**

HTTP code	Response body
200 - OK	<b>VolumeSnapshot</b> schema
201 - Created	<b>VolumeSnapshot</b> schema
202 - Accepted	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty

### 26.2.2.3. /apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots/{name}

Table 26.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeSnapshot
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 26.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a VolumeSnapshot

Table 26.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

Parameter	Type	Description
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 26.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 26.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

## HTTP method

**GET**

**Description**

read the specified VolumeSnapshot

**Table 26.17. Query parameters**

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

**Table 26.18. HTTP responses**

HTTP code	Response body
200 - OK	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified VolumeSnapshot

**Table 26.19. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 26.20. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 26.21. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified VolumeSnapshot

Table 26.22. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 26.23. Body parameters

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshot</b> schema	

Table 26.24. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshot</b> schema
201 - Created	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty



## 26.2.2.4. /apis/snapshot.storage.k8s.io/v1/namespaces/{namespace}/volumesnapshots/{name}

Table 26.25. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeSnapshot
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 26.26. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

read status of the specified VolumeSnapshot

Table 26.27. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 26.28. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty

## HTTP method

**PATCH**

## Description

partially update status of the specified VolumeSnapshot

Table 26.29. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 26.30. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 26.31. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified VolumeSnapshot

**Table 26.32. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 26.33. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshot</b> schema	

Table 26.34. HTTP responses

HTTP code	Reponse body
200 - OK	<b>VolumeSnapshot</b> schema
201 - Created	<b>VolumeSnapshot</b> schema
401 - Unauthorized	Empty

## 26.3. VOLUMESNAPSHOTCLASS [SNAPSHOT.STORAGE.K8S.IO/V1]

### Description

VolumeSnapshotClass specifies parameters that a underlying storage system uses when creating a volume snapshot. A specific VolumeSnapshotClass is used by specifying its name in a VolumeSnapshot object. VolumeSnapshotClasses are non-namespaced

### Type

**object**

### Required

- **deletionPolicy**
- **driver**

### 26.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>deletionPolicy</b>	<b>string</b>	deletionPolicy determines whether a VolumeSnapshotContent created through the VolumeSnapshotClass should be deleted when its bound VolumeSnapshot is deleted. Supported values are "Retain" and "Delete". "Retain" means that the VolumeSnapshotContent and its physical snapshot on underlying storage system are kept. "Delete" means that the VolumeSnapshotContent and its physical snapshot on underlying storage system are deleted. Required.
<b>driver</b>	<b>string</b>	driver is the name of the storage driver that handles this VolumeSnapshotClass. Required.
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>parameters</b>	<b>object (string)</b>	parameters is a key-value map with storage driver specific parameters for creating snapshots. These values are opaque to Kubernetes.

### 26.3.2. API endpoints

The following API endpoints are available:

- **/apis/snapshot.storage.k8s.io/v1/volumesnapshotclasses**
  - **DELETE:** delete collection of VolumeSnapshotClass
  - **GET:** list objects of kind VolumeSnapshotClass
  - **POST:** create a VolumeSnapshotClass
- **/apis/snapshot.storage.k8s.io/v1/volumesnapshotclasses/{name}**
  - **DELETE:** delete a VolumeSnapshotClass
  - **GET:** read the specified VolumeSnapshotClass
  - **PATCH:** partially update the specified VolumeSnapshotClass
  - **PUT:** replace the specified VolumeSnapshotClass

### 26.3.2.1. /apis/snapshot.storage.k8s.io/v1/volumesnapshotclasses

Table 26.35. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of VolumeSnapshotClass

Table 26.36. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 26.37. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

HTTP method

**GET**

**Description**

list objects of kind VolumeSnapshotClass

**Table 26.38. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 26.39. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotClassList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a VolumeSnapshotClass

**Table 26.40. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 26.41. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshotClass</b> schema	

**Table 26.42. HTTP responses**

HTTP code	Response body
200 - OK	<b>VolumeSnapshotClass</b> schema
201 - Created	<b>VolumeSnapshotClass</b> schema
202 - Accepted	<b>VolumeSnapshotClass</b> schema
401 - Unauthorized	Empty

### 26.3.2.2. /apis/snapshot.storage.k8s.io/v1/volumesnapshotclasses/{name}

Table 26.43. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeSnapshotClass

Table 26.44. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a VolumeSnapshotClass

Table 26.45. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 26.46. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 26.47. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified VolumeSnapshotClass

Table 26.48. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 26.49. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotClass</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified VolumeSnapshotClass

Table 26.50. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).



Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 26.51. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 26.52. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotClass</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified VolumeSnapshotClass

Table 26.53. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 26.54. Body parameters

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshotClass</b> schema	

Table 26.55. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotClass</b> schema
201 - Created	<b>VolumeSnapshotClass</b> schema
401 - Unauthorized	Empty

## 26.4. VOLUMESNAPSHOTCONTENT [SNAPSHOT.STORAGE.K8S.IO/V1]

### Description

VolumeSnapshotContent represents the actual "on-disk" snapshot object in the underlying storage system

### Type

**object**

### Required

- **spec**

### 26.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	spec defines properties of a VolumeSnapshotContent created by the underlying storage system. Required.

Property	Type	Description
<b>status</b>	<b>object</b>	status represents the current information of a snapshot.

### 26.4.1.1. .spec

#### Description

spec defines properties of a VolumeSnapshotContent created by the underlying storage system. Required.

#### Type

**object**

#### Required

- **deletionPolicy**
- **driver**
- **source**
- **volumeSnapshotRef**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>deletionPolicy</b>	<b>string</b>	deletionPolicy determines whether this VolumeSnapshotContent and its physical snapshot on the underlying storage system should be deleted when its bound VolumeSnapshot is deleted. Supported values are "Retain" and "Delete". "Retain" means that the VolumeSnapshotContent and its physical snapshot on underlying storage system are kept. "Delete" means that the VolumeSnapshotContent and its physical snapshot on underlying storage system are deleted. For dynamically provisioned snapshots, this field will automatically be filled in by the CSI snapshotter sidecar with the "DeletionPolicy" field defined in the corresponding VolumeSnapshotClass. For pre-existing snapshots, users <b>MUST</b> specify this field when creating the VolumeSnapshotContent object. Required.
<b>driver</b>	<b>string</b>	driver is the name of the CSI driver used to create the physical snapshot on the underlying storage system. This <b>MUST</b> be the same as the name returned by the CSI GetPluginName() call for that driver. Required.
<b>source</b>	<b>object</b>	source specifies whether the snapshot is (or should be) dynamically provisioned or already exists, and just requires a Kubernetes object representation. This field is immutable after creation. Required.

Property	Type	Description
<b>sourceVolumeMode</b>	<b>string</b>	SourceVolumeMode is the mode of the volume whose snapshot is taken. Can be either "Filesystem" or "Block". If not specified, it indicates the source volume's mode is unknown. This field is immutable. This field is an alpha field.
<b>volumeSnapshotClassName</b>	<b>string</b>	name of the VolumeSnapshotClass from which this snapshot was (or will be) created. Note that after provisioning, the VolumeSnapshotClass may be deleted or recreated with different set of values, and as such, should not be referenced post-snapshot creation.
<b>volumeSnapshotRef</b>	<b>object</b>	volumeSnapshotRef specifies the VolumeSnapshot object to which this VolumeSnapshotContent object is bound. VolumeSnapshot.Spec.VolumeSnapshotContentName field must reference to this VolumeSnapshotContent's name for the bidirectional binding to be valid. For a pre-existing VolumeSnapshotContent object, name and namespace of the VolumeSnapshot object MUST be provided for binding to happen. This field is immutable after creation. Required.

### 26.4.1.2. .spec.source

#### Description

source specifies whether the snapshot is (or should be) dynamically provisioned or already exists, and just requires a Kubernetes object representation. This field is immutable after creation. Required.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>snapshotHandle</b>	<b>string</b>	snapshotHandle specifies the CSI "snapshot_id" of a pre-existing snapshot on the underlying storage system for which a Kubernetes object representation was (or should be) created. This field is immutable.
<b>volumeHandle</b>	<b>string</b>	volumeHandle specifies the CSI "volume_id" of the volume from which a snapshot should be dynamically taken from. This field is immutable.

### 26.4.1.3. .spec.volumeSnapshotRef

#### Description

volumeSnapshotRef specifies the VolumeSnapshot object to which this VolumeSnapshotContent object is bound. VolumeSnapshot.Spec.VolumeSnapshotContentName field must reference to this VolumeSnapshotContent's name for the bidirectional binding to be valid. For a pre-existing VolumeSnapshotContent object, name and namespace of the VolumeSnapshot object **MUST** be provided for binding to happen. This field is immutable after creation. Required.

#### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	API version of the referent.

Property	Type	Description
<b>fieldPath</b>	<b>string</b>	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object. TODO: this design is not final and this field is subject to change in the future.
<b>kind</b>	<b>string</b>	Kind of the referent. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>name</b>	<b>string</b>	Name of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names</a>
<b>namespace</b>	<b>string</b>	Namespace of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/">https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/</a>
<b>resourceVersion</b>	<b>string</b>	Specific resourceVersion to which this reference is made, if any. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</a>



Property	Type	Description
<b>uid</b>	<b>string</b>	UID of the referent. More info: <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids">https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids</a>

#### 26.4.1.4. .status

##### Description

status represents the current information of a snapshot.

##### Type

**object**

Property	Type	Description
<b>creationTime</b>	<b>integer</b>	creationTime is the timestamp when the point-in-time snapshot is taken by the underlying storage system. In dynamic snapshot creation case, this field will be filled in by the CSI snapshotter sidecar with the "creation_time" value returned from CSI "CreateSnapshot" gRPC call. For a pre-existing snapshot, this field will be filled with the "creation_time" value returned from the CSI "ListSnapshots" gRPC call if the driver supports it. If not specified, it indicates the creation time is unknown. The format of this field is a Unix nanoseconds time encoded as an int64. On Unix, the command <b>date +%s%N</b> returns the current time in nanoseconds since 1970-01-01 00:00:00 UTC.
<b>error</b>	<b>object</b>	error is the last observed error during snapshot creation, if any. Upon success after retry, this error field will be cleared.

Property	Type	Description
<b>readyToUse</b>	<b>boolean</b>	readyToUse indicates if a snapshot is ready to be used to restore a volume. In dynamic snapshot creation case, this field will be filled in by the CSI snapshotter sidecar with the "ready_to_use" value returned from CSI "CreateSnapshot" gRPC call. For a pre-existing snapshot, this field will be filled with the "ready_to_use" value returned from the CSI "ListSnapshots" gRPC call if the driver supports it, otherwise, this field will be set to "True". If not specified, it means the readiness of a snapshot is unknown.
<b>restoreSize</b>	<b>integer</b>	restoreSize represents the complete size of the snapshot in bytes. In dynamic snapshot creation case, this field will be filled in by the CSI snapshotter sidecar with the "size_bytes" value returned from CSI "CreateSnapshot" gRPC call. For a pre-existing snapshot, this field will be filled with the "size_bytes" value returned from the CSI "ListSnapshots" gRPC call if the driver supports it. When restoring a volume from this snapshot, the size of the volume MUST NOT be smaller than the restoreSize if it is specified, otherwise the restoration will fail. If not specified, it indicates that the size is unknown.
<b>snapshotHandle</b>	<b>string</b>	snapshotHandle is the CSI "snapshot_id" of a snapshot on the underlying storage system. If not specified, it indicates that dynamic snapshot creation has either failed or it is still in progress.

#### 26.4.1.5. .status.error

##### Description

error is the last observed error during snapshot creation, if any. Upon success after retry, this error field will be cleared.

## Type

### object

Property	Type	Description
<b>message</b>	<b>string</b>	message is a string detailing the encountered error during snapshot creation if specified. NOTE: message may be logged, and it should not contain sensitive information.
<b>time</b>	<b>string</b>	time is the timestamp when the error was encountered.

## 26.4.2. API endpoints

The following API endpoints are available:

- **/apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents**
  - **DELETE**: delete collection of VolumeSnapshotContent
  - **GET**: list objects of kind VolumeSnapshotContent
  - **POST**: create a VolumeSnapshotContent
- **/apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents/{name}**
  - **DELETE**: delete a VolumeSnapshotContent
  - **GET**: read the specified VolumeSnapshotContent
  - **PATCH**: partially update the specified VolumeSnapshotContent
  - **PUT**: replace the specified VolumeSnapshotContent
- **/apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents/{name}/status**
  - **GET**: read status of the specified VolumeSnapshotContent
  - **PATCH**: partially update status of the specified VolumeSnapshotContent
  - **PUT**: replace status of the specified VolumeSnapshotContent

### 26.4.2.1. /apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents

Table 26.56. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**DELETE**

## Description

delete collection of VolumeSnapshotContent

Table 26.57. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue`</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 26.58. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

HTTP method

**GET**

**Description**

list objects of kind VolumeSnapshotContent

**Table 26.59. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 26.60. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContentList</b> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a VolumeSnapshotContent

**Table 26.61. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 26.62. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshotContent</b> schema	

**Table 26.63. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">VolumeSnapshotContent</a> schema
201 - Created	<a href="#">VolumeSnapshotContent</a> schema
202 - Accepted	<a href="#">VolumeSnapshotContent</a> schema
401 - Unauthorized	Empty

### 26.4.2.2. /apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents/{name}

Table 26.64. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeSnapshotContent

Table 26.65. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a VolumeSnapshotContent

Table 26.66. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 26.67. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 26.68. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified VolumeSnapshotContent

Table 26.69. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 26.70. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContent</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified VolumeSnapshotContent

Table 26.71. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 26.72. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 26.73. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContent</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified VolumeSnapshotContent

Table 26.74. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 26.75. Body parameters

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshotContent</b> schema	

Table 26.76. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContent</b> schema
201 - Created	<b>VolumeSnapshotContent</b> schema
401 - Unauthorized	Empty

## 26.4.2.3. /apis/snapshot.storage.k8s.io/v1/volumesnapshotcontents/{name}/status

Table 26.77. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeSnapshotContent

Table 26.78. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

read status of the specified VolumeSnapshotContent

Table 26.79. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 26.80. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContent</b> schema
401 - Unauthorized	Empty

## HTTP method

**PATCH**

## Description

partially update status of the specified VolumeSnapshotContent

Table 26.81. Query parameters



Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 26.82. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 26.83. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContent</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified VolumeSnapshotContent

**Table 26.84. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 26.85. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>VolumeSnapshotContent</b> schema	

Table 26.86. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeSnapshotContent</b> schema
201 - Created	<b>VolumeSnapshotContent</b> schema
401 - Unauthorized	Empty

## CHAPTER 27. STORAGE APIS

### 27.1. STORAGE APIS

#### 27.1.1. CSIDriver [storage.k8s.io/v1]

##### Description

CSIDriver captures information about a Container Storage Interface (CSI) volume driver deployed on the cluster. Kubernetes attach detach controller uses this object to determine whether attach is required. Kubelet uses this object to determine whether pod information needs to be passed on mount. CSIDriver objects are non-namespaced.

##### Type

**object**

#### 27.1.2. CSINode [storage.k8s.io/v1]

##### Description

CSINode holds information about all CSI drivers installed on a node. CSI drivers do not need to create the CSINode object directly. As long as they use the node-driver-registrar sidecar container, the kubelet will automatically populate the CSINode object for the CSI driver as part of kubelet plugin registration. CSINode has the same name as a node. If the object is missing, it means either there are no CSI Drivers available on the node, or the Kubelet version is low enough that it doesn't create this object. CSINode has an OwnerReference that points to the corresponding node object.

##### Type

**object**

#### 27.1.3. CSIStorageCapacity [storage.k8s.io/v1]

##### Description

CSIStorageCapacity stores the result of one CSI GetCapacity call. For a given StorageClass, this describes the available capacity in a particular topology segment. This can be used when considering where to instantiate new PersistentVolumes.

For example this can express things like: - StorageClass "standard" has "1234 GiB" available in "topology.kubernetes.io/zone=us-east1" - StorageClass "localssd" has "10 GiB" available in "kubernetes.io/hostname=knnode-abc123"

The following three cases all imply that no capacity is available for a certain combination: - no object exists with suitable topology and storage class name - such an object exists, but the capacity is unset - such an object exists, but the capacity is zero

The producer of these objects can decide which approach is more suitable.

They are consumed by the kube-scheduler when a CSI driver opts into capacity-aware scheduling with CSIDriverSpec.StorageCapacity. The scheduler compares the MaximumVolumeSize against the requested size of pending volumes to filter out unsuitable nodes. If MaximumVolumeSize is unset, it falls back to a comparison against the less precise Capacity. If that is also unset, the scheduler assumes that capacity is insufficient and tries some other node.

##### Type

**object**

### 27.1.4. StorageClass [storage.k8s.io/v1]

#### Description

StorageClass describes the parameters for a class of storage for which PersistentVolumes can be dynamically provisioned.

StorageClasses are non-namespaced; the name of the storage class according to etcd is in ObjectMeta.Name.

#### Type

**object**

### 27.1.5. VolumeAttachment [storage.k8s.io/v1]

#### Description

VolumeAttachment captures the intent to attach or detach the specified volume to/from the specified node.

VolumeAttachment objects are non-namespaced.

#### Type

**object**

## 27.2. CSIDRIVER [STORAGE.K8S.IO/V1]

#### Description

CSIDriver captures information about a Container Storage Interface (CSI) volume driver deployed on the cluster. Kubernetes attach detach controller uses this object to determine whether attach is required. Kubelet uses this object to determine whether pod information needs to be passed on mount. CSIDriver objects are non-namespaced.

#### Type

**object**

#### Required

- **spec**

### 27.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object metadata. metadata.Name indicates the name of the CSI driver that this object refers to; it MUST be the same name returned by the CSI GetPluginName() call for that driver. The driver name must be 63 characters or less, beginning and ending with an alphanumeric character ([a-z0-9A-Z]) with dashes (-), dots (.), and alphanumerics between. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	CSIDriverSpec is the specification of a CSIDriver.

### 27.2.1.1. .spec

#### Description

CSIDriverSpec is the specification of a CSIDriver.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>attachRequired</b>	<b>boolean</b>	<p>attachRequired indicates this CSI volume driver requires an attach operation (because it implements the CSI ControllerPublishVolume() method), and that the Kubernetes attach detach controller should call the attach volume interface which checks the volumeattachment status and waits until the volume is attached before proceeding to mounting. The CSI external-attacher coordinates with CSI volume driver and updates the volumeattachment status when the attach operation is complete. If the CSIDriverRegistry feature gate is enabled and the value is specified to false, the attach operation will be skipped. Otherwise the attach operation will be called.</p> <p>This field is immutable.</p>
<b>fsGroupPolicy</b>	<b>string</b>	<p>fsGroupPolicy defines if the underlying volume supports changing ownership and permission of the volume before being mounted. Refer to the specific FSGroupPolicy values for additional details.</p> <p>This field is immutable.</p> <p>Defaults to ReadWriteOnceWithFSType, which will examine each volume to determine if Kubernetes should modify ownership and permissions of the volume. With the default policy the defined fsGroup will only be applied if a fstype is defined and the volume's access mode contains ReadWriteOnce.</p>

Property	Type	Description
<b>podInfoOnMount</b>	<b>boolean</b>	<p>podInfoOnMount indicates this CSI volume driver requires additional pod information (like podName, podUID, etc.) during mount operations, if set to true. If set to false, pod information will not be passed on mount. Default is false.</p> <p>The CSI driver specifies podInfoOnMount as part of driver deployment. If true, Kubelet will pass pod information as VolumeContext in the CSI NodePublishVolume() calls. The CSI driver is responsible for parsing and validating the information passed in as VolumeContext.</p> <p>The following VolumeContext will be passed if podInfoOnMount is set to true. This list might grow, but the prefix will be used.</p> <p>"csi.storage.k8s.io/pod.name": pod.Name  "csi.storage.k8s.io/pod.namespace": pod.Namespace  "csi.storage.k8s.io/pod.uid": string(pod.UID)  "csi.storage.k8s.io/ephemeral": "true" if the volume is an ephemeral inline volume defined by a CSIVolumeSource, otherwise "false"</p> <p>"csi.storage.k8s.io/ephemeral" is a new feature in Kubernetes 1.16. It is only required for drivers which support both the "Persistent" and "Ephemeral" VolumeLifecycleMode. Other drivers can leave pod info disabled and/or ignore this field. As Kubernetes 1.15 doesn't support this field, drivers can only support one mode when deployed on such a cluster and the deployment determines which mode that is, for example via a command line parameter of the driver.</p> <p>This field is immutable.</p>



Property	Type	Description
<b>requiresRepublish</b>	<b>boolean</b>	<p>requiresRepublish indicates the CSI driver wants <b>NodePublishVolume</b> being periodically called to reflect any possible change in the mounted volume. This field defaults to false.</p> <p>Note: After a successful initial NodePublishVolume call, subsequent calls to NodePublishVolume should only update the contents of the volume. New mount points will not be seen by a running container.</p>
<b>seLinuxMount</b>	<b>boolean</b>	<p>seLinuxMount specifies if the CSI driver supports "-o context" mount option.</p> <p>When "true", the CSI driver must ensure that all volumes provided by this CSI driver can be mounted separately with different <b>-o context</b> options. This is typical for storage backends that provide volumes as filesystems on block devices or as independent shared volumes. Kubernetes will call NodeStage / NodePublish with "-o context=xyz" mount option when mounting a ReadWriteOncePod volume used in Pod that has explicitly set SELinux context. In the future, it may be expanded to other volume AccessModes. In any case, Kubernetes will ensure that the volume is mounted only with a single SELinux context.</p> <p>When "false", Kubernetes won't pass any special SELinux mount options to the driver. This is typical for volumes that represent subdirectories of a bigger shared filesystem.</p> <p>Default is "false".</p>

Property	Type	Description
<b>storageCapacity</b>	<b>boolean</b>	<p>storageCapacity indicates that the CSI volume driver wants pod scheduling to consider the storage capacity that the driver deployment will report by creating CSISStorageCapacity objects with capacity information, if set to true.</p> <p>The check can be enabled immediately when deploying a driver. In that case, provisioning new volumes with late binding will pause until the driver deployment has published some suitable CSISStorageCapacity object.</p> <p>Alternatively, the driver can be deployed with the field unset or false and it can be flipped later when storage capacity information has been published.</p> <p>This field was immutable in Kubernetes <math>\leq</math> 1.22 and now is mutable.</p>
<b>tokenRequests</b>	<b>array</b>	<p>tokenRequests indicates the CSI driver needs pods' service account tokens it is mounting volume for to do necessary authentication. Kubelet will pass the tokens in VolumeContext in the CSI NodePublishVolume calls. The CSI driver should parse and validate the following VolumeContext:</p> <pre>"csi.storage.k8s.io/serviceAccount.tokens": { "&lt;audience&gt;": {   "token": &lt;token&gt;,   "expirationTimestamp":     &lt;expiration timestamp in RFC3339&gt;, }, ... }</pre> <p>Note: Audience in each TokenRequest should be different and at most one token is empty string. To receive a new token after expiry, RequiresRepublish can be used to trigger NodePublishVolume periodically.</p>

Property	Type	Description
<b>tokenRequests[]</b>	<b>object</b>	TokenRequest contains parameters of a service account token.
<b>volumeLifecycleModes</b>	<b>array (string)</b>	<p>volumeLifecycleModes defines what kind of volumes this CSI volume driver supports. The default if the list is empty is "Persistent", which is the usage defined by the CSI specification and implemented in Kubernetes via the usual PV/PVC mechanism.</p> <p>The other mode is "Ephemeral". In this mode, volumes are defined inline inside the pod spec with CSIVolumeSource and their lifecycle is tied to the lifecycle of that pod. A driver has to be aware of this because it is only going to get a NodePublishVolume call for such a volume.</p> <p>For more information about implementing this mode, see <a href="https://kubernetes-csi.github.io/docs/ephemeral-local-volumes.html">https://kubernetes-csi.github.io/docs/ephemeral-local-volumes.html</a> A driver can support one or more of these modes and more modes may be added in the future.</p> <p>This field is beta. This field is immutable.</p>

### 27.2.1.2. .spec.tokenRequests

#### Description

tokenRequests indicates the CSI driver needs pods' service account tokens it is mounting volume for to do necessary authentication. Kubelet will pass the tokens in VolumeContext in the CSI NodePublishVolume calls. The CSI driver should parse and validate the following VolumeContext: "csi.storage.k8s.io/serviceAccount.tokens": { "<audience>": { "token": <token>, "expirationTimestamp": <expiration timestamp in RFC3339>, }, ... }

Note: Audience in each TokenRequest should be different and at most one token is empty string. To receive a new token after expiry, RequiresRepublish can be used to trigger NodePublishVolume periodically.

#### Type

**array**

### 27.2.1.3. .spec.tokenRequests[]

#### Description

TokenRequest contains parameters of a service account token.

#### Type

**object**

#### Required

- **audience**

Property	Type	Description
<b>audience</b>	<b>string</b>	audience is the intended audience of the token in "TokenRequestSpec". It will default to the audiences of kube apiserver.
<b>expirationSeconds</b>	<b>integer</b>	expirationSeconds is the duration of validity of the token in "TokenRequestSpec". It has the same default value of "ExpirationSeconds" in "TokenRequestSpec".

### 27.2.2. API endpoints

The following API endpoints are available:

- **/apis/storage.k8s.io/v1/csidrivers**
  - **DELETE**: delete collection of CSIDriver
  - **GET**: list or watch objects of kind CSIDriver
  - **POST**: create a CSIDriver
- **/apis/storage.k8s.io/v1/watch/csidrivers**
  - **GET**: watch individual changes to a list of CSIDriver. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/storage.k8s.io/v1/csidrivers/{name}**
  - **DELETE**: delete a CSIDriver
  - **GET**: read the specified CSIDriver
  - **PATCH**: partially update the specified CSIDriver
  - **PUT**: replace the specified CSIDriver
- **/apis/storage.k8s.io/v1/watch/csidrivers/{name}**

- **GET**: watch changes to an object of kind CSIDriver. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 27.2.2.1. /apis/storage.k8s.io/v1/csidrivers

Table 27.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of CSIDriver

Table 27.2. Query parameters

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

Parameter	Type	Description
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 27.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.4. HTTP responses



HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind CSIDriver

**Table 27.5. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 27.6. HTTP responses

HTTP code	Reponse body
200 - OK	<b>CSIDriverList</b> schema
401 - Unauthorized	Empty

## HTTP method

**POST****Description**

create a CSIDriver

**Table 27.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 27.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CSIDriver</b> schema	

**Table 27.9. HTTP responses**

HTTP code	Response body
200 - OK	<b>CSIDriver</b> schema

HTTP code	Response body
201 - Created	<b>CSIDriver</b> schema
202 - Accepted	<b>CSIDriver</b> schema
401 - Unauthorized	Empty

### 27.2.2.2. /apis/storage.k8s.io/v1/watch/csidrivers

Table 27.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CSIDriver. deprecated: use the 'watch' parameter with a list operation instead.

Table 27.11. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 27.2.2.3. /apis/storage.k8s.io/v1/csidrivers/{name}

Table 27.12. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CSIDriver

Table 27.13. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a CSIDriver

Table 27.14. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 27.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.16. HTTP responses

HTTP code	Response body
200 - OK	<b>CSIDriver</b> schema
202 - Accepted	<b>CSIDriver</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified CSIDriver

Table 27.17. HTTP responses

HTTP code	Response body
200 - OK	<b>CSIDriver</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified CSIDriver

**Table 27.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 27.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 27.20. HTTP responses

HTTP code	Reponse body
200 - OK	<b>CSIDriver</b> schema
201 - Created	<b>CSIDriver</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified CSIDriver

Table 27.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 27.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CSIDriver</b> schema	

Table 27.23. HTTP responses

HTTP code	Response body
200 - OK	<b>CSIDriver</b> schema
201 - Created	<b>CSIDriver</b> schema
401 - Unauthorized	Empty

#### 27.2.2.4. /apis/storage.k8s.io/v1/watch/csidrivers/{name}

Table 27.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CSIDriver

Table 27.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind CSIDriver. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 27.26. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 27.3. CSINODE [STORAGE.K8S.IO/V1]

### Description

CSINode holds information about all CSI drivers installed on a node. CSI drivers do not need to create the CSINode object directly. As long as they use the node-driver-registrar sidecar container, the kubelet will automatically populate the CSINode object for the CSI driver as part of kubelet plugin registration. CSINode has the same name as a node. If the object is missing, it means either there are no CSI Drivers available on the node, or the Kubelet version is low enough that it doesn't create this object. CSINode has an OwnerReference that points to the corresponding node object.

### Type

**object**

### Required

- **spec**

### 27.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>



Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. metadata.name must be the Kubernetes node name.
<b>spec</b>	<b>object</b>	CSINodeSpec holds information about the specification of all CSI drivers installed on a node

### 27.3.1.1. .spec

#### Description

CSINodeSpec holds information about the specification of all CSI drivers installed on a node

#### Type

**object**

#### Required

- **drivers**

Property	Type	Description
<b>drivers</b>	<b>array</b>	drivers is a list of information of all CSI Drivers existing on a node. If all drivers in the list are uninstalled, this can become empty.
<b>drivers[]</b>	<b>object</b>	CSINodeDriver holds information about the specification of one CSI driver installed on a node

### 27.3.1.2. .spec.drivers

#### Description

drivers is a list of information of all CSI Drivers existing on a node. If all drivers in the list are uninstalled, this can become empty.

#### Type

**array**

### 27.3.1.3. .spec.drivers[]

#### Description

CSINodeDriver holds information about the specification of one CSI driver installed on a node

#### Type

**object****Required**

- **name**
- **nodeID**

Property	Type	Description
<b>allocatable</b>	<b>object</b>	VolumeNodeResources is a set of resource limits for scheduling of volumes.
<b>name</b>	<b>string</b>	name represents the name of the CSI driver that this object refers to. This <b>MUST</b> be the same name returned by the CSI <code>GetPluginName()</code> call for that driver.
<b>nodeID</b>	<b>string</b>	nodeID of the node from the driver point of view. This field enables Kubernetes to communicate with storage systems that do not share the same nomenclature for nodes. For example, Kubernetes may refer to a given node as "node1", but the storage system may refer to the same node as "nodeA". When Kubernetes issues a command to the storage system to attach a volume to a specific node, it can use this field to refer to the node name using the ID that the storage system will understand, e.g. "nodeA" instead of "node1". This field is required.

Property	Type	Description
<b>topologyKeys</b>	<b>array (string)</b>	topologyKeys is the list of keys supported by the driver. When a driver is initialized on a cluster, it provides a set of topology keys that it understands (e.g. "company.com/zone", "company.com/region"). When a driver is initialized on a node, it provides the same topology keys along with values. Kubelet will expose these topology keys as labels on its own node object. When Kubernetes does topology aware provisioning, it can use this list to determine which labels it should retrieve from the node object and pass back to the driver. It is possible for different nodes to use different topology keys. This can be empty if driver does not support topology.

#### 27.3.1.4. .spec.drivers[].allocatable

##### Description

VolumeNodeResources is a set of resource limits for scheduling of volumes.

##### Type

**object**

Property	Type	Description
<b>count</b>	<b>integer</b>	count indicates the maximum number of unique volumes managed by the CSI driver that can be used on a node. A volume that is both attached and mounted on a node is considered to be used once, not twice. The same rule applies for a unique volume that is shared among multiple pods on the same node. If this field is not specified, then the supported number of volumes on this node is unbounded.

#### 27.3.2. API endpoints

The following API endpoints are available:

- **/apis/storage.k8s.io/v1/csinodes**
  - **DELETE**: delete collection of CSINode
  - **GET**: list or watch objects of kind CSINode
  - **POST**: create a CSINode
- **/apis/storage.k8s.io/v1/watch/csinodes**
  - **GET**: watch individual changes to a list of CSINode. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/storage.k8s.io/v1/csinodes/{name}**
  - **DELETE**: delete a CSINode
  - **GET**: read the specified CSINode
  - **PATCH**: partially update the specified CSINode
  - **PUT**: replace the specified CSINode
- **/apis/storage.k8s.io/v1/watch/csinodes/{name}**
  - **GET**: watch changes to an object of kind CSINode. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 27.3.2.1. /apis/storage.k8s.io/v1/csinodes

Table 27.27. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of CSINode

Table 27.28. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>

Parameter	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 27.29. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.30. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind CSINode

Table 27.31. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.



Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 27.32. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CSINodeList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a CSINode

**Table 27.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 27.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CSINode</b> schema	

**Table 27.35. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>CSINode</b> schema

HTTP code	Response body
201 - Created	<b>CSINode</b> schema
202 - Accepted	<b>CSINode</b> schema
401 - Unauthorized	Empty

### 27.3.2.2. /apis/storage.k8s.io/v1/watch/csinodes

Table 27.36. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
GET

**Description**

watch individual changes to a list of CSINode. deprecated: use the 'watch' parameter with a list operation instead.

**Table 27.37. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**27.3.2.3. /apis/storage.k8s.io/v1/csinodes/{name}****Table 27.38. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CSINode

**Table 27.39. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a CSINode

**Table 27.40. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 27.41. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.42. HTTP responses

HTTP code	Response body
200 - OK	<b>CSINode</b> schema
202 - Accepted	<b>CSINode</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified CSINode

Table 27.43. HTTP responses

HTTP code	Response body
200 - OK	<b>CSINode</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified CSINode

**Table 27.44. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 27.45. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 27.46. HTTP responses

HTTP code	Response body
200 - OK	<b>CSINode</b> schema
201 - Created	<b>CSINode</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified CSINode

Table 27.47. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 27.48. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CSINode</b> schema	

Table 27.49. HTTP responses

HTTP code	Response body
200 - OK	<b>CSINode</b> schema
201 - Created	<b>CSINode</b> schema
401 - Unauthorized	Empty

### 27.3.2.4. /apis/storage.k8s.io/v1/watch/csinodes/{name}

Table 27.50. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CSINode

Table 27.51. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind CSINode. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 27.52. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 27.4. CSISTORAGECAPACITY [STORAGE.K8S.IO/V1]

### Description

CSISStorageCapacity stores the result of one CSI GetCapacity call. For a given StorageClass, this describes the available capacity in a particular topology segment. This can be used when considering where to instantiate new PersistentVolumes.

For example this can express things like: - StorageClass "standard" has "1234 GiB" available in "topology.kubernetes.io/zone=us-east1" - StorageClass "localssd" has "10 GiB" available in "kubernetes.io/hostname=knnode-abc123"

The following three cases all imply that no capacity is available for a certain combination: - no object exists with suitable topology and storage class name - such an object exists, but the capacity is unset - such an object exists, but the capacity is zero

The producer of these objects can decide which approach is more suitable.

They are consumed by the kube-scheduler when a CSI driver opts into capacity-aware scheduling with CSIDriverSpec.StorageCapacity. The scheduler compares the MaximumVolumeSize against the requested size of pending volumes to filter out unsuitable nodes. If MaximumVolumeSize is unset, it falls back to a comparison against the less precise Capacity. If that is also unset, the scheduler assumes that capacity is insufficient and tries some other node.

### Type

**object**

### Required

- **storageClassName**

### 27.4.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>



Property	Type	Description
<b>capacity</b>	<b>Quantity</b>	<p>capacity is the value reported by the CSI driver in its GetCapacityResponse for a GetCapacityRequest with topology and parameters that match the previous fields.</p> <p>The semantic is currently (CSI spec 1.2) defined as: The available capacity, in bytes, of the storage that can be used to provision volumes. If not set, that information is currently unavailable.</p>
<b>kind</b>	<b>string</b>	<p>Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a></p>
<b>maximumVolumeSize</b>	<b>Quantity</b>	<p>maximumVolumeSize is the value reported by the CSI driver in its GetCapacityResponse for a GetCapacityRequest with topology and parameters that match the previous fields.</p> <p>This is defined since CSI spec 1.4.0 as the largest size that may be used in a CreateVolumeRequest.capacity_range.required_bytes field to create a volume with the same parameters as those in GetCapacityRequest. The corresponding value in the Kubernetes API is ResourceRequirements.Requests in a volume claim.</p>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	<p>Standard object's metadata. The name has no particular meaning. It must be a DNS subdomain (dots allowed, 253 characters). To ensure that there are no conflicts with other CSI drivers on the cluster, the recommendation is to use <code>csisc-&lt;uuid&gt;</code>, a generated name, or a reverse-domain name which ends with the unique CSI driver name.</p> <p>Objects are namespaced.</p> <p>More info:  <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a></p>
<b>nodeTopology</b>	<b>LabelSelector</b>	<p><code>nodeTopology</code> defines which nodes have access to the storage for which capacity was reported. If not set, the storage is not accessible from any node in the cluster. If empty, the storage is accessible from all nodes. This field is immutable.</p>
<b>storageClassName</b>	<b>string</b>	<p><code>storageClassName</code> represents the name of the <code>StorageClass</code> that the reported capacity applies to. It must meet the same requirements as the name of a <code>StorageClass</code> object (non-empty, DNS subdomain). If that object no longer exists, the <code>CSISStorageCapacity</code> object is obsolete and should be removed by its creator. This field is immutable.</p>

## 27.4.2. API endpoints

The following API endpoints are available:

- **/apis/storage.k8s.io/v1/csistoragecapacities**
  - **GET**: list or watch objects of kind `CSISStorageCapacity`

- **/apis/storage.k8s.io/v1/watch/csistoragecapacities**
  - **GET**: watch individual changes to a list of CSIStorageCapacity. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/storage.k8s.io/v1/namespaces/{namespace}/csistoragecapacities**
  - **DELETE**: delete collection of CSIStorageCapacity
  - **GET**: list or watch objects of kind CSIStorageCapacity
  - **POST**: create a CSIStorageCapacity
- **/apis/storage.k8s.io/v1/watch/namespaces/{namespace}/csistoragecapacities**
  - **GET**: watch individual changes to a list of CSIStorageCapacity. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/storage.k8s.io/v1/namespaces/{namespace}/csistoragecapacities/{name}**
  - **DELETE**: delete a CSIStorageCapacity
  - **GET**: read the specified CSIStorageCapacity
  - **PATCH**: partially update the specified CSIStorageCapacity
  - **PUT**: replace the specified CSIStorageCapacity
- **/apis/storage.k8s.io/v1/watch/namespaces/{namespace}/csistoragecapacities/{name}**
  - **GET**: watch changes to an object of kind CSIStorageCapacity. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 27.4.2.1. /apis/storage.k8s.io/v1/csistoragecapacities

Table 27.53. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

list or watch objects of kind CSISStorageCapacity

**Table 27.54. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>CSISStorageCapacityList</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

### 27.4.2.2. /apis/storage.k8s.io/v1/watch/csistoragecapacities

Table 27.55. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CSISStorageCapacity. deprecated: use the 'watch' parameter with a list operation instead.

**Table 27.56. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 27.4.2.3. /apis/storage.k8s.io/v1/namespaces/{namespace}/csistoragecapacities

Table 27.57. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 27.58. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of CSISStorageCapacity

Table 27.59. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>

Table 27.60. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.61. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind CSIStorageCapacity

Table 27.62. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 27.63. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">CSISStorageCapacityList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a CSISStorageCapacity

**Table 27.64. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 27.65. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>CSISStorageCapacity</b> schema	

**Table 27.66. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>CSISStorageCapacity</b> schema

HTTP code	Response body
201 - Created	<a href="#">CSIStorageCapacity</a> schema
202 - Accepted	<a href="#">CSIStorageCapacity</a> schema
401 - Unauthorized	Empty

#### 27.4.2.4. /apis/storage.k8s.io/v1/watch/namespaces/{namespace}/csistoragecapacities

Table 27.67. Global path parameters

Parameter	Type	Description
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 27.68. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch individual changes to a list of CSISStorageCapacity. deprecated: use the 'watch' parameter with a list operation instead.

Table 27.69. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 27.4.2.5. /apis/storage.k8s.io/v1/namespaces/{namespace}/csistoragecapacities/{name}

Table 27.70. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CSIStorageCapacity
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Table 27.71. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a CSIStorageCapacity

Table 27.72. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 27.73. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.74. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified CSIStorageCapacity

Table 27.75. HTTP responses

HTTP code	Response body
200 - OK	<b>CSIStorageCapacity</b> schema



HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified CSISStorageCapacity

**Table 27.76. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 27.77. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 27.78. HTTP responses

HTTP code	Response body
200 - OK	<b>CSStorageCapacity</b> schema
201 - Created	<b>CSStorageCapacity</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified CSStorageCapacity

Table 27.79. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 27.80. Body parameters

Parameter	Type	Description
<b>body</b>	<b>CSISStorageCapacity</b> schema	

Table 27.81. HTTP responses

HTTP code	Response body
200 - OK	<b>CSISStorageCapacity</b> schema
201 - Created	<b>CSISStorageCapacity</b> schema
401 - Unauthorized	Empty

#### 27.4.2.6. /apis/storage.k8s.io/v1/watch/namespaces/{namespace}/csistoragecapacities/{na

Table 27.82. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the CSISStorageCapacity
<b>namespace</b>	<b>string</b>	object name and auth scope, such as for teams and projects

Parameter	Type	Description
-----------	------	-------------

Table 27.83. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind CSIStorageCapacity. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

**Table 27.84. HTTP responses**

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 27.5. STORAGECLASS [STORAGE.K8S.IO/V1]

### Description

StorageClass describes the parameters for a class of storage for which PersistentVolumes can be dynamically provisioned.

StorageClasses are non-namespaced; the name of the storage class according to etcd is in ObjectMeta.Name.

### Type

**object**

### Required

- **provisioner**

### 27.5.1. Specification

Property	Type	Description
<b>allowVolumeExpansion</b>	<b>boolean</b>	allowVolumeExpansion shows whether the storage class allow volume expand.
<b>allowedTopologies</b>	<b>array</b> <b>(TopologySelectorTerm)</b>	allowedTopologies restrict the node topologies where volumes can be dynamically provisioned. Each volume plugin defines its own supported topology specifications. An empty TopologySelectorTerm list means there is no topology restriction. This field is only honored by servers that enable the VolumeScheduling feature.

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>mountOptions</b>	<b>array (string)</b>	mountOptions controls the mountOptions for dynamically provisioned PersistentVolumes of this storage class. e.g. ["ro", "soft"]. Not validated - mount of the PVs will simply fail if one is invalid.
<b>parameters</b>	<b>object (string)</b>	parameters holds the parameters for the provisioner that should create volumes of this storage class.
<b>provisioner</b>	<b>string</b>	provisioner indicates the type of the provisioner.



Property	Type	Description
<b>reclaimPolicy</b>	<b>string</b>	<p>reclaimPolicy controls the reclaimPolicy for dynamically provisioned PersistentVolumes of this storage class. Defaults to Delete.</p> <p>Possible enum values: - <b>"Delete"</b> means the volume will be deleted from Kubernetes on release from its claim. The volume plugin must support Deletion. - <b>"Recycle"</b> means the volume will be recycled back into the pool of unbound persistent volumes on release from its claim. The volume plugin must support Recycling. - <b>"Retain"</b> means the volume will be left in its current phase (Released) for manual reclamation by the administrator. The default policy is Retain.</p>
<b>volumeBindingMode</b>	<b>string</b>	<p>volumeBindingMode indicates how PersistentVolumeClaims should be provisioned and bound. When unset, VolumeBindingImmediate is used. This field is only honored by servers that enable the VolumeScheduling feature.</p> <p>Possible enum values: - <b>"Immediate"</b> indicates that PersistentVolumeClaims should be immediately provisioned and bound. This is the default mode. - <b>"WaitForFirstConsumer"</b> indicates that PersistentVolumeClaims should not be provisioned and bound until the first Pod is created that references the PersistentVolumeClaim. The volume provisioning and binding will occur during Pod scheduling.</p>

### 27.5.2. API endpoints

The following API endpoints are available:

- `/apis/storage.k8s.io/v1/storageclasses`

- **DELETE**: delete collection of StorageClass
- **GET**: list or watch objects of kind StorageClass
- **POST**: create a StorageClass
- **/apis/storage.k8s.io/v1/watch/storageclasses**
  - **GET**: watch individual changes to a list of StorageClass. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/storage.k8s.io/v1/storageclasses/{name}**
  - **DELETE**: delete a StorageClass
  - **GET**: read the specified StorageClass
  - **PATCH**: partially update the specified StorageClass
  - **PUT**: replace the specified StorageClass
- **/apis/storage.k8s.io/v1/watch/storageclasses/{name}**
  - **GET**: watch changes to an object of kind StorageClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

### 27.5.2.1. /apis/storage.k8s.io/v1/storageclasses

Table 27.85. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of StorageClass

Table 27.86. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>
<b>gracePeriodSeconds</b>	<b>integer</b>	<p>The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.</p>
<b>labelSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their labels. Defaults to everything.</p>

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	<p>Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.</p>
<b>propagationPolicy</b>	<b>string</b>	<p>Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.</p>

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>

Parameter	Type	Description
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 27.87. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.88. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind StorageClass

Table 27.89. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 27.90. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">StorageClassList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST**

**Description**

create a StorageClass

**Table 27.91. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 27.92. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>StorageClass</b> schema	

**Table 27.93. HTTP responses**

HTTP code	Reponse body
200 - OK	<b>StorageClass</b> schema

HTTP code	Response body
201 - Created	<b>StorageClass</b> schema
202 - Accepted	<b>StorageClass</b> schema
401 - Unauthorized	Empty

### 27.5.2.2. /apis/storage.k8s.io/v1/watch/storageclasses

Table 27.94. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method  
GET

## Description

watch individual changes to a list of StorageClass. deprecated: use the 'watch' parameter with a list operation instead.

**Table 27.95. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 27.5.2.3. /apis/storage.k8s.io/v1/storageclasses/{name}

**Table 27.96. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StorageClass

**Table 27.97. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

### DELETE

## Description

delete a StorageClass

**Table 27.98. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 27.99. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.100. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageClass</b> schema
202 - Accepted	<b>StorageClass</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified StorageClass

Table 27.101. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageClass</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified StorageClass

**Table 27.102. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.



Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 27.103. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 27.104. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageClass</b> schema
201 - Created	<b>StorageClass</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified StorageClass

Table 27.105. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 27.106. Body parameters

Parameter	Type	Description
<b>body</b>	<b>StorageClass</b> schema	

Table 27.107. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageClass</b> schema
201 - Created	<b>StorageClass</b> schema
401 - Unauthorized	Empty

#### 27.5.2.4. /apis/storage.k8s.io/v1/watch/storageclasses/{name}

Table 27.108. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StorageClass

Table 27.109. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <code>continue</code> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind StorageClass. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 27.110. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

## 27.6. VOLUMEATTACHMENT [STORAGE.K8S.IO/V1]

### Description

VolumeAttachment captures the intent to attach or detach the specified volume to/from the specified node.

VolumeAttachment objects are non-namespaced.

### Type

**object**

### Required

- **spec**

### 27.6.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	VolumeAttachmentSpec is the specification of a VolumeAttachment request.
<b>status</b>	<b>object</b>	VolumeAttachmentStatus is the status of a VolumeAttachment request.

### 27.6.1.1. .spec

#### Description

VolumeAttachmentSpec is the specification of a VolumeAttachment request.

#### Type

**object**

#### Required

- **attacher**
- **source**
- **nodeName**

Property	Type	Description
<b>attacher</b>	<b>string</b>	attacher indicates the name of the volume driver that MUST handle this request. This is the name returned by GetPluginName().
<b>nodeName</b>	<b>string</b>	nodeName represents the node that the volume should be attached to.

Property	Type	Description
<b>source</b>	<b>object</b>	VolumeAttachmentSource represents a volume that should be attached. Right now only PersistentVolumes can be attached via external attacher, in future we may allow also inline volumes in pods. Exactly one member can be set.

### 27.6.1.2. .spec.source

#### Description

VolumeAttachmentSource represents a volume that should be attached. Right now only PersistentVolumes can be attached via external attacher, in future we may allow also inline volumes in pods. Exactly one member can be set.

#### Type

**object**

Property	Type	Description
<b>inlineVolumeSpec</b>	<b>PersistentVolumeSpec</b>	inlineVolumeSpec contains all the information necessary to attach a persistent volume defined by a pod's inline VolumeSource. This field is populated only for the CSIMigration feature. It contains translated fields from a pod's inline VolumeSource to a PersistentVolumeSpec. This field is beta-level and is only honored by servers that enabled the CSIMigration feature.
<b>persistentVolumeName</b>	<b>string</b>	persistentVolumeName represents the name of the persistent volume to attach.

### 27.6.1.3. .status

#### Description

VolumeAttachmentStatus is the status of a VolumeAttachment request.

#### Type

**object**

#### Required

- **attached**



Property	Type	Description
<b>attachError</b>	<b>object</b>	VolumeError captures an error encountered during a volume operation.
<b>attached</b>	<b>boolean</b>	attached indicates the volume is successfully attached. This field must only be set by the entity completing the attach operation, i.e. the external-attacher.
<b>attachmentMetadata</b>	<b>object (string)</b>	attachmentMetadata is populated with any information returned by the attach operation, upon successful attach, that must be passed into subsequent WaitForAttach or Mount calls. This field must only be set by the entity completing the attach operation, i.e. the external-attacher.
<b>detachError</b>	<b>object</b>	VolumeError captures an error encountered during a volume operation.

#### 27.6.1.4. .status.attachError

##### Description

VolumeError captures an error encountered during a volume operation.

##### Type

**object**

Property	Type	Description
<b>message</b>	<b>string</b>	message represents the error encountered during Attach or Detach operation. This string may be logged, so it should not contain sensitive information.
<b>time</b>	<b>Time</b>	time represents the time the error was encountered.

#### 27.6.1.5. .status.detachError

##### Description

VolumeError captures an error encountered during a volume operation.

Type  
object

Property	Type	Description
<b>message</b>	<b>string</b>	message represents the error encountered during Attach or Detach operation. This string may be logged, so it should not contain sensitive information.
<b>time</b>	<b>Time</b>	time represents the time the error was encountered.

## 27.6.2. API endpoints

The following API endpoints are available:

- **/apis/storage.k8s.io/v1/volumeattachments**
  - **DELETE**: delete collection of VolumeAttachment
  - **GET**: list or watch objects of kind VolumeAttachment
  - **POST**: create a VolumeAttachment
- **/apis/storage.k8s.io/v1/watch/volumeattachments**
  - **GET**: watch individual changes to a list of VolumeAttachment. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/storage.k8s.io/v1/volumeattachments/{name}**
  - **DELETE**: delete a VolumeAttachment
  - **GET**: read the specified VolumeAttachment
  - **PATCH**: partially update the specified VolumeAttachment
  - **PUT**: replace the specified VolumeAttachment
- **/apis/storage.k8s.io/v1/watch/volumeattachments/{name}**
  - **GET**: watch changes to an object of kind VolumeAttachment. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.
- **/apis/storage.k8s.io/v1/volumeattachments/{name}/status**
  - **GET**: read status of the specified VolumeAttachment
  - **PATCH**: partially update status of the specified VolumeAttachment
  - **PUT**: replace status of the specified VolumeAttachment

## 27.6.2.1. /apis/storage.k8s.io/v1/volumeattachments

Table 27.111. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**DELETE**

## Description

delete collection of VolumeAttachment

Table 27.112. Query parameters

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

Parameter	Type	Description
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 27.113. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.114. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind VolumeAttachment

Table 27.115. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 27.116. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">VolumeAttachmentList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a VolumeAttachment

**Table 27.117. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 27.118. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>VolumeAttachment</b> schema	

**Table 27.119. HTTP responses**

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema
201 - Created	<b>VolumeAttachment</b> schema
202 - Accepted	<b>VolumeAttachment</b> schema
401 - Unauthorized	Empty

### 27.6.2.2. /apis/storage.k8s.io/v1/watch/volumeattachments

Table 27.120. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method

**GET****Description**

watch individual changes to a list of VolumeAttachment. deprecated: use the 'watch' parameter with a list operation instead.

Table 27.121. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**27.6.2.3. /apis/storage.k8s.io/v1/volumeattachments/{name}**

Table 27.122. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeAttachment

Table 27.123. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a VolumeAttachment

Table 27.124. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 27.125. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 27.126. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema
202 - Accepted	<b>VolumeAttachment</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified VolumeAttachment

Table 27.127. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified VolumeAttachment

**Table 27.128. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.



Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 27.129. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 27.130. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema
201 - Created	<b>VolumeAttachment</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified VolumeAttachment

Table 27.131. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 27.132. Body parameters

Parameter	Type	Description
<b>body</b>	<b>VolumeAttachment</b> schema	

Table 27.133. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema
201 - Created	<b>VolumeAttachment</b> schema
401 - Unauthorized	Empty

#### 27.6.2.4. /apis/storage.k8s.io/v1/watch/volumeattachments/{name}

Table 27.134. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeAttachment

Table 27.135. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind VolumeAttachment. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 27.136. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 27.6.2.5. /apis/storage.k8s.io/v1/volumeattachments/{name}/status

Table 27.137. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the VolumeAttachment

Table 27.138. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified VolumeAttachment

Table 27.139. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">VolumeAttachment</a> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified VolumeAttachment

Table 27.140. Query parameters

Parameter	Type	Description
-----------	------	-------------

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 27.141. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 27.142. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema
201 - Created	<b>VolumeAttachment</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified VolumeAttachment

**Table 27.143. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 27.144. Body parameters**



Parameter	Type	Description
<b>body</b>	<b>VolumeAttachment</b> schema	

Table 27.145. HTTP responses

HTTP code	Response body
200 - OK	<b>VolumeAttachment</b> schema
201 - Created	<b>VolumeAttachment</b> schema
401 - Unauthorized	Empty

## CHAPTER 28. STORAGE VERSION MIGRATION APIS

### 28.1. STORAGE VERSION MIGRATION APIS

#### 28.1.1. StorageVersionMigration [migration.k8s.io/v1alpha1]

##### Description

StorageVersionMigration represents a migration of stored data to the latest storage version.

##### Type

**object**

### 28.2. STORAGEVERSIONMIGRATION [MIGRATION.K8S.IO/V1ALPHA1]

##### Description

StorageVersionMigration represents a migration of stored data to the latest storage version.

##### Type

**object**

#### 28.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	Specification of the migration.
<b>status</b>	<b>object</b>	Status of the migration.

### 28.2.1.1. .spec

#### Description

Specification of the migration.

#### Type

**object**

#### Required

- **resource**

Property	Type	Description
<b>continueToken</b>	<b>string</b>	The token used in the list options to get the next chunk of objects to migrate. When the <code>.status.conditions</code> indicates the migration is "Running", users can use this token to check the progress of the migration.
<b>resource</b>	<b>object</b>	The resource that is being migrated. The migrator sends requests to the endpoint serving the resource. Immutable.

### 28.2.1.2. .spec.resource

#### Description

The resource that is being migrated. The migrator sends requests to the endpoint serving the resource. Immutable.

#### Type

**object**

Property	Type	Description
<b>group</b>	<b>string</b>	The name of the group.
<b>resource</b>	<b>string</b>	The name of the resource.
<b>version</b>	<b>string</b>	The name of the version.

### 28.2.1.3. .status

#### Description

Status of the migration.

#### Type

**object**

Property	Type	Description
<b>conditions</b>	<b>array</b>	The latest available observations of the migration's current state.
<b>conditions[]</b>	<b>object</b>	Describes the state of a migration at a certain point.

### 28.2.1.4. .status.conditions

#### Description

The latest available observations of the migration's current state.

#### Type

**array**

### 28.2.1.5. .status.conditions[]

#### Description

Describes the state of a migration at a certain point.

#### Type

**object**

#### Required

- **status**
- **type**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>lastUpdateTime</b>	<b>string</b>	The last time this condition was updated.
<b>message</b>	<b>string</b>	A human readable message indicating details about the transition.
<b>reason</b>	<b>string</b>	The reason for the condition's last transition.
<b>status</b>	<b>string</b>	Status of the condition, one of True, False, Unknown.
<b>type</b>	<b>string</b>	Type of the condition.

## 28.2.2. API endpoints

The following API endpoints are available:

- **/apis/migration.k8s.io/v1alpha1/storageversionmigrations**
  - **DELETE**: delete collection of StorageVersionMigration
  - **GET**: list objects of kind StorageVersionMigration
  - **POST**: create a StorageVersionMigration
- **/apis/migration.k8s.io/v1alpha1/storageversionmigrations/{name}**
  - **DELETE**: delete a StorageVersionMigration
  - **GET**: read the specified StorageVersionMigration
  - **PATCH**: partially update the specified StorageVersionMigration
  - **PUT**: replace the specified StorageVersionMigration
- **/apis/migration.k8s.io/v1alpha1/storageversionmigrations/{name}/status**
  - **GET**: read status of the specified StorageVersionMigration
  - **PATCH**: partially update status of the specified StorageVersionMigration
  - **PUT**: replace status of the specified StorageVersionMigration

### 28.2.2.1. /apis/migration.k8s.io/v1alpha1/storageversionmigrations

Table 28.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**DELETE**

## Description

delete collection of StorageVersionMigration

Table 28.2. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 28.3. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

HTTP method



**GET****Description**

list objects of kind StorageVersionMigration

**Table 28.4. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`"k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 28.5. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">StorageVersionMigrationList</a> schema
401 - Unauthorized	Empty

## HTTP method

**POST**

## Description

create a StorageVersionMigration

Table 28.6. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 28.7. Body parameters

Parameter	Type	Description
<b>body</b>	<b>StorageVersionMigration</b> schema	

Table 28.8. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">StorageVersionMigration</a> schema
201 - Created	<a href="#">StorageVersionMigration</a> schema
202 - Accepted	<a href="#">StorageVersionMigration</a> schema
401 - Unauthorized	Empty

### 28.2.2.2. /apis/migration.k8s.io/v1alpha1/storageversionmigrations/{name}

Table 28.9. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StorageVersionMigration

Table 28.10. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### DELETE

#### Description

delete a StorageVersionMigration

Table 28.11. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 28.12. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 28.13. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified StorageVersionMigration

Table 28.14. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 28.15. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageVersionMigration</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified StorageVersionMigration

Table 28.16. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 28.17. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 28.18. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageVersionMigration</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified StorageVersionMigration

Table 28.19. Query parameters



Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 28.20. Body parameters

Parameter	Type	Description
<b>body</b>	<b>StorageVersionMigration</b> schema	

Table 28.21. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageVersionMigration</b> schema
201 - Created	<b>StorageVersionMigration</b> schema
401 - Unauthorized	Empty

### 28.2.2.3. /apis/migration.k8s.io/v1alpha1/storageversionmigrations/{name}/status

Table 28.22. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the StorageVersionMigration

Table 28.23. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

##### GET

#### Description

read status of the specified StorageVersionMigration

Table 28.24. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 28.25. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageVersionMigration</b> schema
401 - Unauthorized	Empty

#### HTTP method

##### PATCH

#### Description

partially update status of the specified StorageVersionMigration

Table 28.26. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 28.27. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 28.28. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageVersionMigration</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified StorageVersionMigration

**Table 28.29. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 28.30. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>StorageVersionMigration</b> schema	

Table 28.31. HTTP responses

HTTP code	Response body
200 - OK	<b>StorageVersionMigration</b> schema
201 - Created	<b>StorageVersionMigration</b> schema
401 - Unauthorized	Empty

## CHAPTER 29. TOPOLVM APIS

### 29.1. TOPOLVM APIS

#### 29.1.1. LogicalVolume [topolvm.io/v1]

##### Description

LogicalVolume is the Schema for the logicalvolumes API

##### Type

**object**

### 29.2. LOGICALVOLUME [TOPOLVM.IO/V1]

##### Description

LogicalVolume is the Schema for the logicalvolumes API

##### Type

**object**

#### 29.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>

Property	Type	Description
<b>metadata</b>	<b>ObjectMeta</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>
<b>spec</b>	<b>object</b>	LogicalVolumeSpec defines the desired state of LogicalVolume
<b>status</b>	<b>object</b>	LogicalVolumeStatus defines the observed state of LogicalVolume

### 29.2.1.1. .spec

#### Description

LogicalVolumeSpec defines the desired state of LogicalVolume

#### Type

**object**

#### Required

- **name**
- **nodeName**
- **size**

Property	Type	Description
<b>accessType</b>	<b>string</b>	'accessType' specifies how the user intends to consume the snapshot logical volume. Set to "ro" when creating a snapshot and to "rw" when restoring a snapshot or creating a clone. This field is populated only when LogicalVolume has a source.
<b>deviceClass</b>	<b>string</b>	
<b>name</b>	<b>string</b>	
<b>nodeName</b>	<b>string</b>	
<b>size</b>	<b>integer-or-string</b>	

Property	Type	Description
<b>source</b>	<b>string</b>	'source' specifies the logicalvolume name of the source; if present. This field is populated only when LogicalVolume has a source.

### 29.2.1.2. .status

#### Description

LogicalVolumeStatus defines the observed state of LogicalVolume

#### Type

**object**

Property	Type	Description
<b>code</b>	<b>integer</b>	A Code is an unsigned 32-bit error code as defined in the gRPC spec.
<b>currentSize</b>	<b>integer-or-string</b>	
<b>message</b>	<b>string</b>	
<b>volumeID</b>	<b>string</b>	INSERT ADDITIONAL STATUS FIELD - define observed state of cluster Important: Run "make" to regenerate code after modifying this file

### 29.2.2. API endpoints

The following API endpoints are available:

- **/apis/topolvm.io/v1/logicalvolumes**
  - **DELETE**: delete collection of LogicalVolume
  - **GET**: list objects of kind LogicalVolume
  - **POST**: create a LogicalVolume
- **/apis/topolvm.io/v1/logicalvolumes/{name}**
  - **DELETE**: delete a LogicalVolume
  - **GET**: read the specified LogicalVolume



- **PATCH**: partially update the specified LogicalVolume
- **PUT**: replace the specified LogicalVolume
- **/apis/topolvm.io/v1/logicalvolumes/{name}/status**
  - **GET**: read status of the specified LogicalVolume
  - **PATCH**: partially update status of the specified LogicalVolume
  - **PUT**: replace status of the specified LogicalVolume

### 29.2.2.1. /apis/topolvm.io/v1/logicalvolumes

Table 29.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete collection of LogicalVolume

Table 29.2. Query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 29.3. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
401 - Unauthorized	Empty

HTTP method

**GET**

**Description**

list objects of kind LogicalVolume

**Table 29.4. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 29.5. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">LogicalVolumeList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a LogicalVolume

**Table 29.6. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 29.7. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>LogicalVolume</b> schema	

**Table 29.8. HTTP responses**



HTTP code	Response body
200 - OK	<b>LogicalVolume</b> schema
201 - Created	<b>LogicalVolume</b> schema
202 - Accepted	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

### 29.2.2.2. /apis/topolvm.io/v1/logicalvolumes/{name}

Table 29.9. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the LogicalVolume

Table 29.10. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

#### HTTP method

#### **DELETE**

#### Description

delete a LogicalVolume

Table 29.11. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 29.12. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 29.13. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified LogicalVolume

Table 29.14. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 29.15. HTTP responses

HTTP code	Response body
200 - OK	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified LogicalVolume

Table 29.16. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 29.17. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 29.18. HTTP responses

HTTP code	Response body
200 - OK	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified LogicalVolume

Table 29.19. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Table 29.20. Body parameters

Parameter	Type	Description
<b>body</b>	<b>LogicalVolume</b> schema	

Table 29.21. HTTP responses

HTTP code	Response body
200 - OK	<b>LogicalVolume</b> schema
201 - Created	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

## 29.2.2.3. /apis/topolvm.io/v1/logicalvolumes/{name}/status

Table 29.22. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the LogicalVolume

Table 29.23. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

**GET**

## Description

read status of the specified LogicalVolume

Table 29.24. Query parameters

Parameter	Type	Description
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Table 29.25. HTTP responses

HTTP code	Response body
200 - OK	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

## HTTP method

**PATCH**

## Description

partially update status of the specified LogicalVolume

Table 29.26. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 29.27. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 29.28. HTTP responses

HTTP code	Response body
200 - OK	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace status of the specified LogicalVolume

**Table 29.29. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 29.30. Body parameters**



Parameter	Type	Description
<b>body</b>	<b>LogicalVolume</b> schema	

Table 29.31. HTTP responses

HTTP code	Reponse body
200 - OK	<b>LogicalVolume</b> schema
201 - Created	<b>LogicalVolume</b> schema
401 - Unauthorized	Empty

## CHAPTER 30. WEBHOOK APIS

### 30.1. WEBHOOK APIS

#### 30.1.1. MutatingWebhookConfiguration [admissionregistration.k8s.io/v1]

##### Description

MutatingWebhookConfiguration describes the configuration of and admission webhook that accept or reject and may change the object.

##### Type

**object**

#### 30.1.2. ValidatingWebhookConfiguration [admissionregistration.k8s.io/v1]

##### Description

ValidatingWebhookConfiguration describes the configuration of and admission webhook that accept or reject and object without changing it.

##### Type

**object**

### 30.2. MUTATINGWEBHOOKCONFIGURATION [ADMISSIONREGISTRATION.K8S.IO/V1]

##### Description

MutatingWebhookConfiguration describes the configuration of and admission webhook that accept or reject and may change the object.

##### Type

**object**

#### 30.2.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>

Property	Type	Description
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object metadata; More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a> .
<b>webhooks</b>	<b>array</b>	Webhooks is a list of webhooks and the affected resources and operations.
<b>webhooks[]</b>	<b>object</b>	MutatingWebhook describes an admission webhook and the resources and operations it applies to.

### 30.2.1.1. .webhooks

#### Description

Webhooks is a list of webhooks and the affected resources and operations.

#### Type

**array**

### 30.2.1.2. .webhooks[]

#### Description

MutatingWebhook describes an admission webhook and the resources and operations it applies to.

#### Type

**object**

#### Required

- **name**
- **clientConfig**
- **sideEffects**

- **admissionReviewVersions**

Property	Type	Description
<b>admissionReviewVersions</b>	<b>array (string)</b>	AdmissionReviewVersions is an ordered list of preferred <b>AdmissionReview</b> versions the Webhook expects. API server will try to use first version in the list which it supports. If none of the versions specified in this list supported by API server, validation will fail for this object. If a persisted webhook configuration specifies allowed versions and does not include any versions known to the API Server, calls to the webhook will fail and be subject to the failure policy.
<b>clientConfig</b>	<b>object</b>	WebhookClientConfig contains the information to make a TLS connection with the webhook
<b>failurePolicy</b>	<b>string</b>	<p>FailurePolicy defines how unrecognized errors from the admission endpoint are handled - allowed values are Ignore or Fail. Defaults to Fail.</p> <p>Possible enum values: - <b>"Fail"</b> means that an error calling the webhook causes the admission to fail. - <b>"Ignore"</b> means that an error calling the webhook is ignored.</p>

Property	Type	Description
<b>matchConditions</b>	<b>array</b>	<p>MatchConditions is a list of conditions that must be met for a request to be sent to this webhook. Match conditions filter requests that have already been matched by the rules, namespaceSelector, and objectSelector. An empty list of matchConditions matches all requests. There are a maximum of 64 match conditions allowed.</p> <p>The exact matching logic is (in order): 1. If ANY matchCondition evaluates to FALSE, the webhook is skipped. 2. If ALL matchConditions evaluate to TRUE, the webhook is called. 3. If any matchCondition evaluates to an error (but none are FALSE): - If failurePolicy=Fail, reject the request - If failurePolicy=Ignore, the error is ignored and the webhook is skipped</p> <p>This is an alpha feature and managed by the AdmissionWebhookMatchConditions feature gate.</p>
<b>matchConditions[]</b>	<b>object</b>	MatchCondition represents a condition which must be fulfilled for a request to be sent to a webhook.

Property	Type	Description
<b>matchPolicy</b>	<b>string</b>	<p>matchPolicy defines how the "rules" list is used to match incoming requests. Allowed values are "Exact" or "Equivalent".</p> <ul style="list-style-type: none"> <li>- Exact: match a request only if it exactly matches a specified rule. For example, if deployments can be modified via apps/v1, apps/v1beta1, and extensions/v1beta1, but "rules" only included <b>apiGroups: ["apps"], apiVersions:["v1"], resources: ["deployments"]</b>, a request to apps/v1beta1 or extensions/v1beta1 would not be sent to the webhook.</li> <li>- Equivalent: match a request if modifies a resource listed in rules, even via another API group or version. For example, if deployments can be modified via apps/v1, apps/v1beta1, and extensions/v1beta1, and "rules" only included <b>apiGroups: ["apps"], apiVersions:["v1"], resources: ["deployments"]</b>, a request to apps/v1beta1 or extensions/v1beta1 would be converted to apps/v1 and sent to the webhook.</li> </ul> <p>Defaults to "Equivalent"</p> <p>Possible enum values: - <b>"Equivalent"</b> means requests should be sent to the webhook if they modify a resource listed in rules via another API group or version. - <b>"Exact"</b> means requests should only be sent to the webhook if they exactly match a given rule.</p>

Property	Type	Description
<b>name</b>	<b>string</b>	The name of the admission webhook. Name should be fully qualified, e.g., imagepolicy.kubernetes.io, where "imagepolicy" is the name of the webhook, and kubernetes.io is the name of the organization. Required.
<b>namespaceSelector</b>	<b>LabelSelector</b>	<p>NamespaceSelector decides whether to run the webhook on an object based on whether the namespace for that object matches the selector. If the object itself is a namespace, the matching is performed on object.metadata.labels. If the object is another cluster scoped resource, it never skips the webhook.</p> <p>For example, to run the webhook on any objects whose namespace is not associated with "runlevel" of "0" or "1"; you will set the selector as follows: "namespaceSelector": { "matchExpressions": [ { "key": "runlevel", "operator": "NotIn", "values": [ "0", "1" ] } ] }</p> <p>If instead you want to only run the webhook on any objects whose namespace is associated with the "environment" of "prod" or "staging"; you will set the selector as follows: "namespaceSelector": { "matchExpressions": [ { "key": "environment", "operator": "In", "values": [ "prod", "staging" ] } ] }</p> <p>See <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/</a> for more examples of label selectors.</p> <p>Default to the empty LabelSelector, which matches everything.</p>

Property	Type	Description
<b>objectSelector</b>	<b>LabelSelector</b>	ObjectSelector decides whether to run the webhook based on if the object has matching labels. objectSelector is evaluated against both the oldObject and newObject that would be sent to the webhook, and is considered to match if either object matches the selector. A null object (oldObject in the case of create, or newObject in the case of delete) or an object that cannot have labels (like a DeploymentRollback or a PodProxyOptions object) is not considered to match. Use the object selector only if the webhook is opt-in, because end users may skip the admission webhook by setting the labels. Default to the empty LabelSelector, which matches everything.



Property	Type	Description
<b>reinvocationPolicy</b>	<b>string</b>	<p>reinvocationPolicy indicates whether this webhook should be called multiple times as part of a single admission evaluation. Allowed values are "Never" and "IfNeeded".</p> <p>Never: the webhook will not be called more than once in a single admission evaluation.</p> <p>IfNeeded: the webhook will be called at least one additional time as part of the admission evaluation if the object being admitted is modified by other admission plugins after the initial webhook call. Webhooks that specify this option <b>must</b> be idempotent, able to process objects they previously admitted.</p> <p>Note: * the number of additional invocations is not guaranteed to be exactly one. * if additional invocations result in further modifications to the object, webhooks are not guaranteed to be invoked again. * webhooks that use this option may be reordered to minimize the number of additional invocations. * to validate an object after all mutations are guaranteed complete, use a validating admission webhook instead.</p> <p>Defaults to "Never".</p> <p>Possible enum values: - <b>"IfNeeded"</b> indicates that the webhook may be called at least one additional time as part of the admission evaluation if the object being admitted is modified by other admission plugins after the initial webhook call. - <b>"Never"</b> indicates that the webhook must not be called more than once in a single admission evaluation.</p>

Property	Type	Description
<b>rules</b>	<b>array</b>	Rules describes what operations on what resources/subresources the webhook cares about. The webhook cares about an operation if it matches <i>any</i> Rule. However, in order to prevent ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks from putting the cluster in a state which cannot be recovered from without completely disabling the plugin, ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks are never called on admission requests for ValidatingWebhookConfiguration and MutatingWebhookConfiguration objects.
<b>rules[]</b>	<b>object</b>	RuleWithOperations is a tuple of Operations and Resources. It is recommended to make sure that all the tuple expansions are valid.

Property	Type	Description
<b>sideEffects</b>	<b>string</b>	<p>SideEffects states whether this webhook has side effects. Acceptable values are: None, NoneOnDryRun (webhooks created via v1beta1 may also specify Some or Unknown). Webhooks with side effects MUST implement a reconciliation system, since a request may be rejected by a future step in the admission chain and the side effects therefore need to be undone. Requests with the dryRun attribute will be auto-rejected if they match a webhook with sideEffects == Unknown or Some.</p> <p>Possible enum values: - <b>"None"</b> means that calling the webhook will have no side effects. - <b>"NoneOnDryRun"</b> means that calling the webhook will possibly have side effects, but if the request being reviewed has the dry-run attribute, the side effects will be suppressed. - <b>"Some"</b> means that calling the webhook will possibly have side effects. If a request with the dry-run attribute would trigger a call to this webhook, the request will instead fail. - <b>"Unknown"</b> means that no information is known about the side effects of calling the webhook. If a request with the dry-run attribute would trigger a call to this webhook, the request will instead fail.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>TimeoutSeconds specifies the timeout for this webhook. After the timeout passes, the webhook call will be ignored or the API call will fail based on the failure policy. The timeout value must be between 1 and 30 seconds. Default to 10 seconds.</p>

### 30.2.1.3. .webhooks[].clientConfig

#### Description

WebhookClientConfig contains the information to make a TLS connection with the webhook

Type

**object**

Property	Type	Description
<b>caBundle</b>	<b>string</b>	<b>caBundle</b> is a PEM encoded CA bundle which will be used to validate the webhook's server certificate. If unspecified, system trust roots on the apiserver are used.
<b>service</b>	<b>object</b>	ServiceReference holds a reference to Service.legacy.k8s.io

Property	Type	Description
<b>url</b>	<b>string</b>	<p><b>url</b> gives the location of the webhook, in standard URL form (<b>scheme://host:port/path</b>). Exactly one of <b>url</b> or <b>service</b> must be specified.</p> <p>The <b>host</b> should not refer to a service running in the cluster; use the <b>service</b> field instead. The host might be resolved via external DNS in some apiservers (e.g., <b>kube-apiserver</b> cannot resolve in-cluster DNS as that would be a layering violation). <b>host</b> may also be an IP address.</p> <p>Please note that using <b>localhost</b> or <b>127.0.0.1</b> as a <b>host</b> is risky unless you take great care to run this webhook on all hosts which run an apiserver which might need to make calls to this webhook. Such installs are likely to be non-portable, i.e., not easy to turn up in a new cluster.</p> <p>The scheme must be "https"; the URL must begin with "https://".</p> <p>A path is optional, and if present may be any string permissible in a URL. You may use the path to pass an arbitrary string to the webhook, for example, a cluster identifier.</p> <p>Attempting to use a user or basic auth e.g. "user:password@" is not allowed. Fragments ("#...") and query parameters ("?...") are not allowed, either.</p>

#### 30.2.1.4. .webhooks[].clientConfig.service

##### Description

ServiceReference holds a reference to Service.legacy.k8s.io

##### Type

**object**

**Required**

- **namespace**
- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	<b>name</b> is the name of the service. Required
<b>namespace</b>	<b>string</b>	<b>namespace</b> is the namespace of the service. Required
<b>path</b>	<b>string</b>	<b>path</b> is an optional URL path which will be sent in any request to this service.
<b>port</b>	<b>integer</b>	If specified, the port on the service that hosting webhook. Default to 443 for backward compatibility. <b>port</b> should be a valid port number (1-65535, inclusive).

**30.2.1.5. .webhooks[].matchConditions****Description**

MatchConditions is a list of conditions that must be met for a request to be sent to this webhook. Match conditions filter requests that have already been matched by the rules, namespaceSelector, and objectSelector. An empty list of matchConditions matches all requests. There are a maximum of 64 match conditions allowed.

The exact matching logic is (in order): 1. If ANY matchCondition evaluates to FALSE, the webhook is skipped. 2. If ALL matchConditions evaluate to TRUE, the webhook is called. 3. If any matchCondition evaluates to an error (but none are FALSE): - If failurePolicy=Fail, reject the request - If failurePolicy=Ignore, the error is ignored and the webhook is skipped

This is an alpha feature and managed by the AdmissionWebhookMatchConditions feature gate.

**Type**

**array**

**30.2.1.6. .webhooks[].matchConditions[]****Description**

MatchCondition represents a condition which must be fulfilled for a request to be sent to a webhook.

**Type**

**object**

**Required**

- **name**

- **name**
- **expression**

Property	Type	Description
<b>expression</b>	<b>string</b>	<p>Expression represents the expression which will be evaluated by CEL. Must evaluate to bool. CEL expressions have access to the contents of the AdmissionRequest and Authorizer, organized into CEL variables:</p> <p>'object' - The object from the incoming request. The value is null for DELETE requests.</p> <p>'oldObject' - The existing object. The value is null for CREATE requests.</p> <p>'request' - Attributes of the admission request(/pkg/apis/admission/types.go#AdmissionRequest).</p> <p>'authorizer' - A CEL Authorizer. May be used to perform authorization checks for the principal (user or service account) of the request. See <a href="https://pkg.go.dev/k8s.io/apiserver/pkg/cel/library#Authorizer">https://pkg.go.dev/k8s.io/apiserver/pkg/cel/library#Authorizer</a></p> <p>'authorizer.requestResource' - A CEL ResourceCheck constructed from the 'authorizer' and configured with the request resource. Documentation on CEL: <a href="https://kubernetes.io/docs/reference/using-api/cel/">https://kubernetes.io/docs/reference/using-api/cel/</a></p> <p>Required.</p>

Property	Type	Description
<b>name</b>	<b>string</b>	Name is an identifier for this match condition, used for strategic merging of MatchConditions, as well as providing an identifier for logging purposes. A good name should be descriptive of the associated expression. Name must be a qualified name consisting of alphanumeric characters, '-', '_' or '.', and must start and end with an alphanumeric character (e.g. 'MyName', or 'my.name', or '123-abc', regex used for validation is '([A-Za-z0-9][-A-Za-z0-9.]*)?[A-Za-z0-9]!') with an optional DNS subdomain prefix and '/' (e.g. 'example.com/MyName')
		Required.

### 30.2.1.7. .webhooks[].rules

#### Description

Rules describes what operations on what resources/subresources the webhook cares about. The webhook cares about an operation if it matches *any* Rule. However, in order to prevent ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks from putting the cluster in a state which cannot be recovered from without completely disabling the plugin, ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks are never called on admission requests for ValidatingWebhookConfiguration and MutatingWebhookConfiguration objects.

#### Type

**array**

### 30.2.1.8. .webhooks[].rules[]

#### Description

RuleWithOperations is a tuple of Operations and Resources. It is recommended to make sure that all the tuple expansions are valid.

#### Type

**object**

Property	Type	Description
<b>apiGroups</b>	<b>array (string)</b>	APIGroups is the API groups the resources belong to. " <b>is all groups. If "</b> is present, the length of the slice must be one. Required.



Property	Type	Description
<b>apiVersions</b>	<b>array (string)</b>	APIVersions is the API versions the resources belong to. <b>" is all versions. If "</b> is present, the length of the slice must be one. Required.
<b>operations</b>	<b>array (string)</b>	Operations is the operations the admission hook cares about - CREATE, UPDATE, DELETE, CONNECT or * for all of those operations and any future admission operations that are added. If '*' is present, the length of the slice must be one. Required.
<b>resources</b>	<b>array (string)</b>	<p>Resources is a list of resources this rule applies to.</p> <p>For example: 'pods' means pods. 'pods/log' means the log subresource of pods. <b>" means all resources, but not subresources. 'pods/' means all subresources of pods. '/scale' means all scale subresources. '/'* means all resources and their subresources.</b></p> <p>If wildcard is present, the validation rule will ensure resources do not overlap with each other.</p> <p>Depending on the enclosing object, subresources might not be allowed. Required.</p>

Property	Type	Description
<b>scope</b>	<b>string</b>	scope specifies the scope of this rule. Valid values are "Cluster", "Namespaced", and "" "Cluster" means that only cluster-scoped resources will match this rule. Namespace API objects are cluster-scoped. "Namespaced" means that only namespaced resources will match this rule. "" means that there are no scope restrictions. Subresources match the scope of their parent resource. Default is "*".

### 30.2.2. API endpoints

The following API endpoints are available:

- **/apis/admissionregistration.k8s.io/v1/mutatingwebhookconfigurations**
  - **DELETE:** delete collection of MutatingWebhookConfiguration
  - **GET:** list or watch objects of kind MutatingWebhookConfiguration
  - **POST:** create a MutatingWebhookConfiguration
- **/apis/admissionregistration.k8s.io/v1/watch/mutatingwebhookconfigurations**
  - **GET:** watch individual changes to a list of MutatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/admissionregistration.k8s.io/v1/mutatingwebhookconfigurations/{name}**
  - **DELETE:** delete a MutatingWebhookConfiguration
  - **GET:** read the specified MutatingWebhookConfiguration
  - **PATCH:** partially update the specified MutatingWebhookConfiguration
  - **PUT:** replace the specified MutatingWebhookConfiguration
- **/apis/admissionregistration.k8s.io/v1/watch/mutatingwebhookconfigurations/{name}**
  - **GET:** watch changes to an object of kind MutatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 30.2.2.1. /apis/admissionregistration.k8s.io/v1/mutatingwebhookconfigurations

Table 30.1. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

## HTTP method

### DELETE

#### Description

delete collection of MutatingWebhookConfiguration

Table 30.2. Query parameters

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	<p>When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed</p>
<b>fieldSelector</b>	<b>string</b>	<p>A selector to restrict the list of returned objects by their fields. Defaults to everything.</p>

Parameter	Type	Description
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

Parameter	Type	Description
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 30.3. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 30.4. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind MutatingWebhookConfiguration

**Table 30.5. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>



Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 30.6. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">MutatingWebhookConfigurationList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a MutatingWebhookConfiguration

**Table 30.7. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 30.8. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>MutatingWebhookConfiguration</b> schema	

**Table 30.9. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">MutatingWebhookConfiguration</a> schema
201 - Created	<a href="#">MutatingWebhookConfiguration</a> schema
202 - Accepted	<a href="#">MutatingWebhookConfiguration</a> schema
401 - Unauthorized	Empty

### 30.2.2.2. /apis/admissionregistration.k8s.io/v1/watch/mutatingwebhookconfigurations

Table 30.10. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method

**GET****Description**

watch individual changes to a list of MutatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead.

**Table 30.11. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**30.2.2.3. /apis/admissionregistration.k8s.io/v1/mutatingwebhookconfigurations/{name}****Table 30.12. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the MutatingWebhookConfiguration

**Table 30.13. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a MutatingWebhookConfiguration

**Table 30.14. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 30.15. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 30.16. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified MutatingWebhookConfiguration

Table 30.17. HTTP responses

HTTP code	Response body
200 - OK	<b>MutatingWebhookConfiguration</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified MutatingWebhookConfiguration

**Table 30.18. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.



Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 30.19. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 30.20. HTTP responses

HTTP code	Response body
200 - OK	<b>MutatingWebhookConfiguration</b> schema
201 - Created	<b>MutatingWebhookConfiguration</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified MutatingWebhookConfiguration

Table 30.21. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 30.22. Body parameters

Parameter	Type	Description
<b>body</b>	<b>MutatingWebhookConfiguration</b> schema	

Table 30.23. HTTP responses

HTTP code	Response body
200 - OK	<b>MutatingWebhookConfiguration</b> schema
201 - Created	<b>MutatingWebhookConfiguration</b> schema
401 - Unauthorized	Empty

#### 30.2.2.4. /apis/admissionregistration.k8s.io/v1/watch/mutatingwebhookconfigurations/{name}

Table 30.24. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the MutatingWebhookConfiguration

Table 30.25. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind MutatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 30.26. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

### 30.3. VALIDATINGWEBHOOKCONFIGURATION [ADMISSIONREGISTRATION.K8S.IO/V1]

#### Description

ValidatingWebhookConfiguration describes the configuration of and admission webhook that accept or reject and object without changing it.

#### Type

**object**

#### 30.3.1. Specification

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>ObjectMeta</b>	Standard object metadata; More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a> .

Property	Type	Description
<b>webhooks</b>	<b>array</b>	Webhooks is a list of webhooks and the affected resources and operations.
<b>webhooks[]</b>	<b>object</b>	ValidatingWebhook describes an admission webhook and the resources and operations it applies to.

### 30.3.1.1. .webhooks

#### Description

Webhooks is a list of webhooks and the affected resources and operations.

#### Type

**array**

### 30.3.1.2. .webhooks[]

#### Description

ValidatingWebhook describes an admission webhook and the resources and operations it applies to.

#### Type

**object**

#### Required

- **name**
- **clientConfig**
- **sideEffects**
- **admissionReviewVersions**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>admissionReviewVersions</b>	<b>array (string)</b>	AdmissionReviewVersions is an ordered list of preferred <b>AdmissionReview</b> versions the Webhook expects. API server will try to use first version in the list which it supports. If none of the versions specified in this list supported by API server, validation will fail for this object. If a persisted webhook configuration specifies allowed versions and does not include any versions known to the API Server, calls to the webhook will fail and be subject to the failure policy.
<b>clientConfig</b>	<b>object</b>	WebhookClientConfig contains the information to make a TLS connection with the webhook
<b>failurePolicy</b>	<b>string</b>	FailurePolicy defines how unrecognized errors from the admission endpoint are handled - allowed values are Ignore or Fail. Defaults to Fail.  Possible enum values: - <b>"Fail"</b> means that an error calling the webhook causes the admission to fail. - <b>"Ignore"</b> means that an error calling the webhook is ignored.



Property	Type	Description
<b>matchConditions</b>	<b>array</b>	<p>MatchConditions is a list of conditions that must be met for a request to be sent to this webhook. Match conditions filter requests that have already been matched by the rules, namespaceSelector, and objectSelector. An empty list of matchConditions matches all requests. There are a maximum of 64 match conditions allowed.</p> <p>The exact matching logic is (in order): 1. If ANY matchCondition evaluates to FALSE, the webhook is skipped. 2. If ALL matchConditions evaluate to TRUE, the webhook is called. 3. If any matchCondition evaluates to an error (but none are FALSE): - If failurePolicy=Fail, reject the request - If failurePolicy=Ignore, the error is ignored and the webhook is skipped</p> <p>This is an alpha feature and managed by the AdmissionWebhookMatchConditions feature gate.</p>
<b>matchConditions[]</b>	<b>object</b>	MatchCondition represents a condition which must be fulfilled for a request to be sent to a webhook.

Property	Type	Description
<b>matchPolicy</b>	<b>string</b>	<p>matchPolicy defines how the "rules" list is used to match incoming requests. Allowed values are "Exact" or "Equivalent".</p> <ul style="list-style-type: none"> <li>- Exact: match a request only if it exactly matches a specified rule. For example, if deployments can be modified via apps/v1, apps/v1beta1, and extensions/v1beta1, but "rules" only included <b>apiGroups: ["apps"], apiVersions: ["v1"], resources: ["deployments"]</b>, a request to apps/v1beta1 or extensions/v1beta1 would not be sent to the webhook.</li> <li>- Equivalent: match a request if modifies a resource listed in rules, even via another API group or version. For example, if deployments can be modified via apps/v1, apps/v1beta1, and extensions/v1beta1, and "rules" only included <b>apiGroups: ["apps"], apiVersions: ["v1"], resources: ["deployments"]</b>, a request to apps/v1beta1 or extensions/v1beta1 would be converted to apps/v1 and sent to the webhook.</li> </ul> <p>Defaults to "Equivalent"</p> <p>Possible enum values: - <b>"Equivalent"</b> means requests should be sent to the webhook if they modify a resource listed in rules via another API group or version. - <b>"Exact"</b> means requests should only be sent to the webhook if they exactly match a given rule.</p>
<b>name</b>	<b>string</b>	<p>The name of the admission webhook. Name should be fully qualified, e.g., imagepolicy.kubernetes.io, where "imagepolicy" is the name of the webhook, and kubernetes.io is the name of the organization. Required.</p>

Property	Type	Description
<b>namespaceSelector</b>	<b>LabelSelector</b>	<p>NamespaceSelector decides whether to run the webhook on an object based on whether the namespace for that object matches the selector. If the object itself is a namespace, the matching is performed on object.metadata.labels. If the object is another cluster scoped resource, it never skips the webhook.</p> <p>For example, to run the webhook on any objects whose namespace is not associated with "runlevel" of "0" or "1"; you will set the selector as follows: "namespaceSelector": { "matchExpressions": [ { "key": "runlevel", "operator": "NotIn", "values": [ "0", "1" ] } ] }</p> <p>If instead you want to only run the webhook on any objects whose namespace is associated with the "environment" of "prod" or "staging"; you will set the selector as follows: "namespaceSelector": { "matchExpressions": [ { "key": "environment", "operator": "In", "values": [ "prod", "staging" ] } ] }</p> <p>See <a href="https://kubernetes.io/docs/concepts/overview/working-with-objects/labels">https://kubernetes.io/docs/concepts/overview/working-with-objects/labels</a> for more examples of label selectors.</p> <p>Default to the empty LabelSelector, which matches everything.</p>

Property	Type	Description
<b>objectSelector</b>	<b>LabelSelector</b>	ObjectSelector decides whether to run the webhook based on if the object has matching labels. objectSelector is evaluated against both the oldObject and newObject that would be sent to the webhook, and is considered to match if either object matches the selector. A null object (oldObject in the case of create, or newObject in the case of delete) or an object that cannot have labels (like a DeploymentRollback or a PodProxyOptions object) is not considered to match. Use the object selector only if the webhook is opt-in, because end users may skip the admission webhook by setting the labels. Default to the empty LabelSelector, which matches everything.
<b>rules</b>	<b>array</b>	Rules describes what operations on what resources/subresources the webhook cares about. The webhook cares about an operation if it matches <i>any</i> Rule. However, in order to prevent ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks from putting the cluster in a state which cannot be recovered from without completely disabling the plugin, ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks are never called on admission requests for ValidatingWebhookConfiguration and MutatingWebhookConfiguration objects.
<b>rules[]</b>	<b>object</b>	RuleWithOperations is a tuple of Operations and Resources. It is recommended to make sure that all the tuple expansions are valid.

Property	Type	Description
<b>sideEffects</b>	<b>string</b>	<p>SideEffects states whether this webhook has side effects. Acceptable values are: None, NoneOnDryRun (webhooks created via v1beta1 may also specify Some or Unknown). Webhooks with side effects MUST implement a reconciliation system, since a request may be rejected by a future step in the admission chain and the side effects therefore need to be undone. Requests with the dryRun attribute will be auto-rejected if they match a webhook with sideEffects == Unknown or Some.</p> <p>Possible enum values: - <b>"None"</b> means that calling the webhook will have no side effects. - <b>"NoneOnDryRun"</b> means that calling the webhook will possibly have side effects, but if the request being reviewed has the dry-run attribute, the side effects will be suppressed. - <b>"Some"</b> means that calling the webhook will possibly have side effects. If a request with the dry-run attribute would trigger a call to this webhook, the request will instead fail. - <b>"Unknown"</b> means that no information is known about the side effects of calling the webhook. If a request with the dry-run attribute would trigger a call to this webhook, the request will instead fail.</p>
<b>timeoutSeconds</b>	<b>integer</b>	<p>TimeoutSeconds specifies the timeout for this webhook. After the timeout passes, the webhook call will be ignored or the API call will fail based on the failure policy. The timeout value must be between 1 and 30 seconds. Default to 10 seconds.</p>

### 30.3.1.3. .webhooks[].clientConfig

#### Description

WebhookClientConfig contains the information to make a TLS connection with the webhook

Type

**object**

Property	Type	Description
<b>caBundle</b>	<b>string</b>	<b>caBundle</b> is a PEM encoded CA bundle which will be used to validate the webhook's server certificate. If unspecified, system trust roots on the apiserver are used.
<b>service</b>	<b>object</b>	ServiceReference holds a reference to Service.legacy.k8s.io

Property	Type	Description
<b>url</b>	<b>string</b>	<p><b>url</b> gives the location of the webhook, in standard URL form (<b>scheme://host:port/path</b>). Exactly one of <b>url</b> or <b>service</b> must be specified.</p> <p>The <b>host</b> should not refer to a service running in the cluster; use the <b>service</b> field instead. The host might be resolved via external DNS in some apiservers (e.g., <b>kube-apiserver</b> cannot resolve in-cluster DNS as that would be a layering violation). <b>host</b> may also be an IP address.</p> <p>Please note that using <b>localhost</b> or <b>127.0.0.1</b> as a <b>host</b> is risky unless you take great care to run this webhook on all hosts which run an apiserver which might need to make calls to this webhook. Such installs are likely to be non-portable, i.e., not easy to turn up in a new cluster.</p> <p>The scheme must be "https"; the URL must begin with "https://".</p> <p>A path is optional, and if present may be any string permissible in a URL. You may use the path to pass an arbitrary string to the webhook, for example, a cluster identifier.</p> <p>Attempting to use a user or basic auth e.g. "user:password@" is not allowed. Fragments ("#...") and query parameters ("?...") are not allowed, either.</p>

#### 30.3.1.4. .webhooks[].clientConfig.service

##### Description

ServiceReference holds a reference to Service.legacy.k8s.io

##### Type

**object**

**Required**

- **namespace**
- **name**

Property	Type	Description
<b>name</b>	<b>string</b>	<b>name</b> is the name of the service. Required
<b>namespace</b>	<b>string</b>	<b>namespace</b> is the namespace of the service. Required
<b>path</b>	<b>string</b>	<b>path</b> is an optional URL path which will be sent in any request to this service.
<b>port</b>	<b>integer</b>	If specified, the port on the service that hosting webhook. Default to 443 for backward compatibility. <b>port</b> should be a valid port number (1-65535, inclusive).

**30.3.1.5. .webhooks[].matchConditions****Description**

MatchConditions is a list of conditions that must be met for a request to be sent to this webhook. Match conditions filter requests that have already been matched by the rules, namespaceSelector, and objectSelector. An empty list of matchConditions matches all requests. There are a maximum of 64 match conditions allowed.

The exact matching logic is (in order): 1. If ANY matchCondition evaluates to FALSE, the webhook is skipped. 2. If ALL matchConditions evaluate to TRUE, the webhook is called. 3. If any matchCondition evaluates to an error (but none are FALSE): - If failurePolicy=Fail, reject the request - If failurePolicy=Ignore, the error is ignored and the webhook is skipped

This is an alpha feature and managed by the AdmissionWebhookMatchConditions feature gate.

**Type**

**array**

**30.3.1.6. .webhooks[].matchConditions[]****Description**

MatchCondition represents a condition which must be fulfilled for a request to be sent to a webhook.

**Type**

**object**

**Required**

- **name**



- **name**
- **expression**

Property	Type	Description
<b>expression</b>	<b>string</b>	<p>Expression represents the expression which will be evaluated by CEL. Must evaluate to bool. CEL expressions have access to the contents of the AdmissionRequest and Authorizer, organized into CEL variables:</p> <p>'object' - The object from the incoming request. The value is null for DELETE requests.</p> <p>'oldObject' - The existing object. The value is null for CREATE requests.</p> <p>'request' - Attributes of the admission request (/pkg/apis/admission/types.go#AdmissionRequest).</p> <p>'authorizer' - A CEL Authorizer. May be used to perform authorization checks for the principal (user or service account) of the request. See <a href="https://pkg.go.dev/k8s.io/apiserver/pkg/cel/library#Authorizer">https://pkg.go.dev/k8s.io/apiserver/pkg/cel/library#Authorizer</a></p> <p>'authorizer.requestResource' - A CEL ResourceCheck constructed from the 'authorizer' and configured with the request resource. Documentation on CEL: <a href="https://kubernetes.io/docs/reference/using-api/cel/">https://kubernetes.io/docs/reference/using-api/cel/</a></p> <p>Required.</p>

Property	Type	Description
<b>name</b>	<b>string</b>	Name is an identifier for this match condition, used for strategic merging of MatchConditions, as well as providing an identifier for logging purposes. A good name should be descriptive of the associated expression. Name must be a qualified name consisting of alphanumeric characters, '-', '_' or '.', and must start and end with an alphanumeric character (e.g. 'MyName', or 'my.name', or '123-abc', regex used for validation is '([A-Za-z0-9](- _ \.)*)?[A-Za-z0-9]!') with an optional DNS subdomain prefix and '/' (e.g. 'example.com/MyName')
		Required.

### 30.3.1.7. .webhooks[].rules

#### Description

Rules describes what operations on what resources/subresources the webhook cares about. The webhook cares about an operation if it matches *any* Rule. However, in order to prevent ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks from putting the cluster in a state which cannot be recovered from without completely disabling the plugin, ValidatingAdmissionWebhooks and MutatingAdmissionWebhooks are never called on admission requests for ValidatingWebhookConfiguration and MutatingWebhookConfiguration objects.

#### Type

**array**

### 30.3.1.8. .webhooks[].rules[]

#### Description

RuleWithOperations is a tuple of Operations and Resources. It is recommended to make sure that all the tuple expansions are valid.

#### Type

**object**

Property	Type	Description
<b>apiGroups</b>	<b>array (string)</b>	APIGroups is the API groups the resources belong to. " <b>is all groups. If "</b> is present, the length of the slice must be one. Required.

Property	Type	Description
<b>apiVersions</b>	<b>array (string)</b>	APIVersions is the API versions the resources belong to. <b>" is all versions. If "</b> is present, the length of the slice must be one. Required.
<b>operations</b>	<b>array (string)</b>	Operations is the operations the admission hook cares about - CREATE, UPDATE, DELETE, CONNECT or * for all of those operations and any future admission operations that are added. If '*' is present, the length of the slice must be one. Required.
<b>resources</b>	<b>array (string)</b>	<p>Resources is a list of resources this rule applies to.</p> <p>For example: 'pods' means pods. 'pods/log' means the log subresource of pods. <b>" means all resources, but not subresources. 'pods/' means all subresources of pods. '/scale' means all scale subresources. '/'* means all resources and their subresources.</b></p> <p>If wildcard is present, the validation rule will ensure resources do not overlap with each other.</p> <p>Depending on the enclosing object, subresources might not be allowed. Required.</p>

Property	Type	Description
<b>scope</b>	<b>string</b>	scope specifies the scope of this rule. Valid values are "Cluster", "Namespaced", and "" "Cluster" means that only cluster-scoped resources will match this rule. Namespace API objects are cluster-scoped. "Namespaced" means that only namespaced resources will match this rule. "" means that there are no scope restrictions. Subresources match the scope of their parent resource. Default is "*".

### 30.3.2. API endpoints

The following API endpoints are available:

- **/apis/admissionregistration.k8s.io/v1/validatingwebhookconfigurations**
  - **DELETE:** delete collection of ValidatingWebhookConfiguration
  - **GET:** list or watch objects of kind ValidatingWebhookConfiguration
  - **POST:** create a ValidatingWebhookConfiguration
- **/apis/admissionregistration.k8s.io/v1/watch/validatingwebhookconfigurations**
  - **GET:** watch individual changes to a list of ValidatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead.
- **/apis/admissionregistration.k8s.io/v1/validatingwebhookconfigurations/{name}**
  - **DELETE:** delete a ValidatingWebhookConfiguration
  - **GET:** read the specified ValidatingWebhookConfiguration
  - **PATCH:** partially update the specified ValidatingWebhookConfiguration
  - **PUT:** replace the specified ValidatingWebhookConfiguration
- **/apis/admissionregistration.k8s.io/v1/watch/validatingwebhookconfigurations/{name}**
  - **GET:** watch changes to an object of kind ValidatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

#### 30.3.2.1. /apis/admissionregistration.k8s.io/v1/validatingwebhookconfigurations

Table 30.27. Global query parameters

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

HTTP method

**DELETE**

Description

delete collection of ValidatingWebhookConfiguration

Table 30.28. Query parameters

Parameter	Type	Description
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

Parameter	Type	Description
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
<b>resourceVersion</b>	<b>string</b>	resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset
<b>resourceVersionMatch</b>	<b>string</b>	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.  Defaults to unset

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end`: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

Table 30.29. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 30.30. HTTP responses

HTTP code	Reponse body
200 - OK	<b>Status</b> schema



HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****GET****Description**

list or watch objects of kind ValidatingWebhookConfiguration

**Table 30.31. Query parameters**

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.

Parameter	Type	Description
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end": "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

Table 30.32. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ValidatingWebhookConfigurationList</a> schema
401 - Unauthorized	Empty

HTTP method

**POST****Description**

create a ValidatingWebhookConfiguration

**Table 30.33. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

**Table 30.34. Body parameters**

Parameter	Type	Description
<b>body</b>	<b>ValidatingWebhookConfiguration</b> schema	

**Table 30.35. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">ValidatingWebhookConfiguration</a> schema
201 - Created	<a href="#">ValidatingWebhookConfiguration</a> schema
202 - Accepted	<a href="#">ValidatingWebhookConfiguration</a> schema
401 - Unauthorized	Empty

### 30.3.2.2. /apis/admissionregistration.k8s.io/v1/watch/validatingwebhookconfigurations

Table 30.36. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

Parameter	Type	Description
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>sendInitialEvents=true</code> may be set together with <code>watch=true</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>k8s.io/initial-events-end: "true"</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>sendInitialEvents</code> option is set, we require <code>resourceVersionMatch</code> option to also be set. The semantic of the watch request is as following: - <code>resourceVersionMatch = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>resourceVersion</code>" and the bookmark event is send when the state is synced to a <code>resourceVersion</code> at least as fresh as the one provided by the ListOptions. If <code>resourceVersion</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>resourceVersionMatch</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>resourceVersion=""</code> or <code>resourceVersion="0"</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

HTTP method

**GET****Description**

watch individual changes to a list of ValidatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead.

**Table 30.37. HTTP responses**

HTTP code	Response body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty

**30.3.2.3. /apis/admissionregistration.k8s.io/v1/validatingwebhookconfigurations/{name}****Table 30.38. Global path parameters**

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ValidatingWebhookConfiguration

**Table 30.39. Global query parameters**

Parameter	Type	Description
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.

**HTTP method****DELETE****Description**

delete a ValidatingWebhookConfiguration

**Table 30.40. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>gracePeriodSeconds</b>	<b>integer</b>	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.



Parameter	Type	Description
<b>orphanDependents</b>	<b>boolean</b>	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
<b>propagationPolicy</b>	<b>string</b>	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

Table 30.41. Body parameters

Parameter	Type	Description
<b>body</b>	<b>DeleteOptions</b> schema	

Table 30.42. HTTP responses

HTTP code	Response body
200 - OK	<b>Status</b> schema
202 - Accepted	<b>Status</b> schema
401 - Unauthorized	Empty

**HTTP method****GET****Description**

read the specified ValidatingWebhookConfiguration

Table 30.43. HTTP responses

HTTP code	Response body
200 - OK	<b>ValidatingWebhookConfiguration</b> schema

HTTP code	Response body
401 - Unauthorized	Empty

**HTTP method****PATCH****Description**

partially update the specified ValidatingWebhookConfiguration

**Table 30.44. Query parameters**

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: - Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23. - Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+ - Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.

Parameter	Type	Description
<b>force</b>	<b>boolean</b>	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.

Table 30.45. Body parameters

Parameter	Type	Description
<b>body</b>	<b>Patch</b> schema	

Table 30.46. HTTP responses

HTTP code	Reponse body
200 - OK	<b>ValidatingWebhookConfiguration</b> schema
201 - Created	<b>ValidatingWebhookConfiguration</b> schema
401 - Unauthorized	Empty

**HTTP method****PUT****Description**

replace the specified ValidatingWebhookConfiguration

Table 30.47. Query parameters

Parameter	Type	Description
<b>dryRun</b>	<b>string</b>	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
<b>fieldManager</b>	<b>string</b>	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by <a href="https://golang.org/pkg/unicode/#IsPrint">https://golang.org/pkg/unicode/#IsPrint</a> .

Parameter	Type	Description
<b>fieldValidation</b>	<b>string</b>	fieldValidation instructs the server on how to handle objects in the request (POST/PUT/PATCH) containing unknown or duplicate fields. Valid values are: <ul style="list-style-type: none"> <li>- Ignore: This will ignore any unknown fields that are silently dropped from the object, and will ignore all but the last duplicate field that the decoder encounters. This is the default behavior prior to v1.23.</li> <li>- Warn: This will send a warning via the standard warning response header for each unknown field that is dropped from the object, and for each duplicate field that is encountered. The request will still succeed if there are no other errors, and will only persist the last of any duplicate fields. This is the default in v1.23+</li> <li>- Strict: This will fail the request with a BadRequest error if any unknown fields would be dropped from the object, or if any duplicate fields are present. The error returned from the server will contain all unknown and duplicate fields encountered.</li> </ul>

Table 30.48. Body parameters

Parameter	Type	Description
<b>body</b>	<a href="#">ValidatingWebhookConfiguration</a> schema	

Table 30.49. HTTP responses

HTTP code	Response body
200 - OK	<a href="#">ValidatingWebhookConfiguration</a> schema
201 - Created	<a href="#">ValidatingWebhookConfiguration</a> schema
401 - Unauthorized	Empty

### 30.3.2.4. /apis/admissionregistration.k8s.io/v1/watch/validatingwebhookconfigurations/{name}

Table 30.50. Global path parameters

Parameter	Type	Description
<b>name</b>	<b>string</b>	name of the ValidatingWebhookConfiguration

Table 30.51. Global query parameters

Parameter	Type	Description
<b>allowWatchBookmarks</b>	<b>boolean</b>	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored.
<b>continue</b>	<b>string</b>	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
<b>fieldSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their fields. Defaults to everything.
<b>labelSelector</b>	<b>string</b>	A selector to restrict the list of returned objects by their labels. Defaults to everything.

Parameter	Type	Description
<b>limit</b>	<b>integer</b>	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the `continue` field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
<b>pretty</b>	<b>string</b>	If 'true', then the output is pretty printed.
<b>resourceVersion</b>	<b>string</b>	<p>resourceVersion sets a constraint on what resource versions a request may be served from. See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>
<b>resourceVersionMatch</b>	<b>string</b>	<p>resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See <a href="https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions">https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions</a> for details.</p> <p>Defaults to unset</p>

Parameter	Type	Description
<b>sendInitialEvents</b>	<b>boolean</b>	<p><code>`sendInitialEvents=true`</code> may be set together with <code>`watch=true`</code>. In that case, the watch stream will begin with synthetic events to produce the current state of objects in the collection. Once all such events have been sent, a synthetic "Bookmark" event will be sent. The bookmark will report the ResourceVersion (RV) corresponding to the set of objects, and be marked with <code>`k8s.io/initial-events-end: "true"`</code> annotation. Afterwards, the watch stream will proceed as usual, sending watch events corresponding to changes (subsequent to the RV) to objects watched.</p> <p>When <code>`sendInitialEvents`</code> option is set, we require <code>`resourceVersionMatch`</code> option to also be set. The semantic of the watch request is as following: - <code>`resourceVersionMatch` = NotOlderThan</code> is interpreted as "data at least as new as the provided <code>`resourceVersion`</code>" and the bookmark event is send when the state is synced to a <code>`resourceVersion`</code> at least as fresh as the one provided by the ListOptions. If <code>`resourceVersion`</code> is unset, this is interpreted as "consistent read" and the bookmark event is send when the state is synced at least to the moment when request started being processed. - <code>`resourceVersionMatch`</code> set to any other value or unset Invalid error is returned.</p> <p>Defaults to true if <code>`resourceVersion=""`</code> or <code>`resourceVersion="0"`</code> (for backward compatibility reasons) and to false otherwise.</p>
<b>timeoutSeconds</b>	<b>integer</b>	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
<b>watch</b>	<b>boolean</b>	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

**HTTP method****GET****Description**

watch changes to an object of kind ValidatingWebhookConfiguration. deprecated: use the 'watch' parameter with a list operation instead, filtered to a single item with the 'fieldSelector' parameter.

Table 30.52. HTTP responses

HTTP code	Reponse body
200 - OK	<a href="#">WatchEvent</a> schema
401 - Unauthorized	Empty